Essays in Public Finance and Taxation

by

Paul Raymond Organ

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Business and Economics)
in the University of Michigan
2022

Doctoral Committee:

Professor Joel Slemrod, Chair
Assistant Professor Ashley Craig
Assistant Professor Edward Fox
Professor James R. Hines, Jr.

Paul Raymond Organ

prorgan@umich.edu

ORCID iD: 0000-0002-6805-3438

© Paul R. Organ 2022

Dedication

I dedicate this dissertation to the friendship and memory of Jim Pagels. Jim was a caring and loyal friend, immensely curious, and a deep and careful thinker. I hope that these essays, and my future work, help carry forward Jim's enthusiasm and commitment to making the world a better place.

Acknowledgements

I am indebted to many people for their support of me and this dissertation.

Thank you to my family for your consistent love and care: my fiancée, Sarah Elizabeth; my parents, Anne and Jim; and my brothers and sisters-in-law; Andrew and Sarah, and Daniel and Bridget. Thanks also to my nephew, Harry, and niece, Georgia, for reminding me that there are at least a few great things in life that are not directly tax-related.

Thank you also to my dissertation committee: Joel Slemrod, Jim Hines, Ash Craig, and Ed Fox. Over the years you have been a steady and reliable source of support, guidance, and inspiration. My research has benefitted immensely from your feedback, as has my overall development as a public finance economist.

Many others at Michigan also deserve my thanks. To the BE faculty at Ross, thank you for your willingness to attend and offer helpful feedback at my brown bag lunches, even though I know you might have been thinking, "Really, another talk about taxes?" To my colleagues, thank you for your friendship, collaboration, and, at times, commiseration: Tyler, Aaron, Evelyn, Jaedo, Mike, Ani, Nate, Emir, Emily, Owen, Tobi, and many others. Thank you to the UM econ trivia team for being a productive distraction, especially Chris, Melissa, Yeliz, and Levi. Thank you to the students who came before me and laid a path for me to follow: Ellen, Andrew, Will, and Max; and to Mary Ceccanese for building a welcoming Michigan PF community. And to Linda and Joe, thank you for giving me a place to stay when that is just what I needed.

Those outside of Michigan also deserve my thanks. To my co-authors, thank you for working with me and putting up with my sometimes too-frequent emails: Alex Turk, Alex Ruda, Brian Galle, Allen Prohofsky, and Chad Angaretis. Many thanks to those at the IRS who supported my research there: John, Pat, Alicia, Anne, Tom, and others. Thanks to Jeff Hoopes for helpful advice. To my professors at GW, thank you for starting me on the path towards this PhD, especially Tony Yezer, Brian Boulier, and Joann Weiner. To my mentors and former colleagues, thank you for helping me develop as a person and as a professional: Elaine Buckberg, Seth Carpenter, George Oldfield, Ben Sacks, Fabricio Nuñez, Mark Berkman, and many others. And, finally, to my friends, Sam, Lucas, Ben, Anna, Ethan, Ally, Nick, Kaushik, Colette, Jacob, Sarah, Ashley, Cady, Mitch, Nate, and others, thank you for your continued support.

Table of Contents

List of Tables

List of Figures

Abstract

This dissertation analyzes individual behavioral responses to tax policy and administration. All three chapters use administrative tax microdata to assess the intended and unintended consequences of policies whose goals are to increase tax compliance or improve tax enforcement and debt collection. Together, these essays expand our understanding of how individuals respond to the tax system, and provide empirical findings that can inform future improvements to that system.

The first chapter explores a recent increase in U.S. citizenship renunciation and its connection to tax rules and enforcement. The U.S. tax system applies to its citizens' worldwide incomes and estates, whether those citizens live in the U.S. or abroad. Fully escaping the U.S. tax system requires dropping U.S. citizenship, and in recent years a growing number of individuals have done so. I use administrative tax microdata to answer three questions: Who is renouncing U.S. citizenship? Why are they renouncing? What are the policy consequences? I show that the recent increase is mainly driven by those who have for many years lived abroad, rather than by individuals leaving the U.S., and that these renunciations are primarily a response to increased compliance costs, not tax liabilities. I also present evidence of a strong response to a net worth threshold affecting the cost of renunciation. I conclude by showing that most of those renouncing had no or little pre-renunciation U.S. tax liability, suggesting their renunciations are likely to have minimal revenue impact. Overall, the evidence shows that the tax system does affect individuals' citizenship decisions, though renunciations are still relatively uncommon.

The second chapter, co-authored with Alex Ruda, Joel Slemrod, and Alex Turk, studies the IRS' passport certification and revocation process. Traditional penalties for tax noncompliance are financial, but many jurisdictions now also use non-monetary tools, including collateral sanctions that deny access to some government-provided service. To learn about the effectiveness of one such penalty, we examine a U.S. policy restricting passport access for taxpayers with substantial tax debt, known as "certification." We take advantage of an RCT during the policy rollout and find small but positive effects on taxpayer compliance of the certification notice sent to all eligible taxpayers. We then study a subset of certified taxpayers who were denied a passport-related request and find an immediate and strong positive effect of the denial on compliance actions.

The third chapter, co-authored with Chad Angaretis, Brian Galle, and Allen Prohofsky, studies California's Top 500 tax delinquent publication program. Many U.S. states and countries around the world publicly disclose tax debtors to encourage compliance. Little is known about the effectiveness of these programs. Using administrative tax microdata from California's "Top 500" disclosure program, we study whether notices of imminent publication affect payment and other compliance outcomes, as well as whether these notices affect subsequent reported earnings. We estimate the direct effect of the letter sent to the 500 highest-balance, publication-eligible taxpayers to be additional revenue of between $2.8 and $7.2 million annually, with no evidence of an impact on subsequent reported earnings. We also estimate an upper bound on the deadweight loss caused by publication of non-compliers, and conclude that the program generates positive net social welfare. Together, these results suggest that delinquent taxpayer disclosure can be an efficient tax enforcement tool, at least among the relatively high-income population we study.

Chapter 1. Citizenship and Taxes: Evaluating the Effects of the U.S. Tax System on Individuals' Citizenship Decisions[1]

1.1. Introduction and motivation

The number of individuals renouncing U.S. citizenship has risen sharply in the last decade, from roughly 500 a year in the early 2000s to more than 4,000 each year from 2013-2018. This increase has been noted in academic literature, with some suggesting a potential connection to the U.S. tax system (e.g., Kudrle 2015; De Simone, Lester, and Markle 2020). Record-high renunciations also appear frequently in the public press, with articles often describing those renouncing as wealthy or high-income and referring to their dropping of citizenship as an attempt to flee a tax system they deem too burdensome. The connection between the tax system and citizenship has important implications for the design of U.S. tax policy. If the recent increase in citizenship renunciation is mainly a response to tax burdens by wealthy and high-income U.S. taxpayers, this would constrain the degree of progressivity that is feasible under the existing tax system. Conversely, if the recent increase in renunciation has other explanations, policymakers have more flexibility in setting progressivity.

[1] For helpful comments, suggestions, and support I thank my dissertation committee: Joel Slemrod, Jim Hines, Ash Craig, and Ed Fox; my thanks also go to Katarzyna Bilicka, Sebastien Bradley, Dhammika Dharmapala, Gabe Ehrlich, Jeff Hoopes, Daniel Reck, Max Risch, Molly Saunders-Scott, Bill Strang, and seminar participants at the University of Michigan, the IRS, the U.S. Treasury Office of Tax Analysis, the 2020 National Tax Association Annual Conference, the 2021 IIPF Annual Congress, and the 2021 Oxford CBT Doctoral Conference. I am especially grateful to John Guyton, Anne Herlache, Thomas Hertz, Pat Langetieg, Alicia Miller, Annette Portz, Alex Turk, and Carlos Zepeda at the IRS for their support of this work. All data work for this project involving confidential taxpayer information was done on IRS computers by IRS employees, and at no time was confidential taxpayer data ever outside of the IRS computing environment. The author is a Student Volunteer with the IRS. The views and opinions presented in this paper reflect those of the author and do not necessarily reflect the views or the official position of the Internal Revenue Service. All results have been reviewed to ensure that no confidential information is disclosed.

In this paper, I use novel administrative tax microdata to provide a more complete and nuanced description of those renouncing citizenship than was previously possible. I show that there is indeed a connection between the tax system and citizenship renunciation, but that most of the recent increase in renunciations is explained by increasing costs of tax compliance for U.S. citizens already living abroad, rather than wealthy or high-income taxpayers leaving the U.S. Further, most of those renouncing citizenship had no or little pre-renunciation U.S. tax liability (net of foreign tax credits and income exclusions), reinforcing that renunciations should mainly be thought of as a response to compliance costs, rather than to tax liabilities.

The U.S. is one of a handful of countries that tax their citizens' worldwide income and estates.[2] As a result, policymakers have frequently raised concerns about U.S. citizens dropping citizenship to reduce their tax burden. The first legislation intended to discourage tax-motivated expatriation was passed in the 1960s. Several high-profile departures in the 1990s prompted new laws related to citizenship renunciation. Since 1998 the names of those dropping U.S. citizenship have been published in the Federal Register. Further substantial changes to the expatriation tax system[3] were passed in 2004 and 2008, including the introduction of a mark-to-market exit tax for certain individuals. Since 2008, the relevant changes have been in tax enforcement, starting with legal actions targeting tax evasion by U.S. citizens in Switzerland, and leading to a broader increase in offshore financial enforcement under the Foreign Account Tax Compliance Act (FATCA). Enacted in 2010, FATCA introduced new requirements for Foreign Financial Institutions (FFIs) to report on their U.S. citizen clients, and was intended to discourage tax non-compliance by U.S.

[2] Only two other countries, Eritrea and Myanmar, similarly tax their citizens regardless of residence. Eritrea levies a flat income tax of 2% on its citizens living abroad; Myanmar applies the same rates to its citizens' income, whether derived at home or abroad.

[3] I use the term "expatriation tax system" to refer to the laws and tax regulations which govern expatriation and citizenship renunciation; these include filing and reporting requirements, and tax liabilities incurred at and after renunciation. Following previous literature and the terminology of related legislation, I use the term "expatriation" to mean giving up U.S. citizenship, rather than merely moving abroad.

taxpayers using foreign accounts. Despite, or perhaps because of, these efforts, citizenship renunciations have continued, with annual counts rising markedly in recent years. Between 2005 and 2018, more than 35,000 individuals with at least $48 billion of combined reported net worth renounced their U.S. citizenship.

Who exactly is renouncing U.S. citizenship? Using administrative tax microdata, I provide detailed information about the population of individuals dropping U.S. citizenship from 1998 to 2018. The recent increase in renunciations has come mainly from those who have long filed U.S. taxes from abroad (and thus likely lived abroad), rather than from individuals who lived in the U.S. choosing to move abroad. Those renouncing citizenship are on average higher-income and wealthier than the U.S. population. More than one-third of those renouncing and reporting net worth reported having more than $1 million, compared with estimates of 5-10% for all U.S. individuals. The number of ultra-wealthy, however, is small. Just 3% of those with reported net worth had more than $10 million; this share is higher than the corresponding U.S. population estimate of 1%, but still represents only about 500 individuals total who renounced from 2005 to 2018. Renunciation is concentrated in relatively few destination jurisdictions, with the top five (Canada, Switzerland, the United Kingdom, Germany, and Hong Kong) accounting for more than half of the total. These facts provide new information about those renouncing citizenship; prior studies have been limited to publicly available information, which in practice has meant only the quarterly counts compiled from the names of renouncers published in the Federal Register.

Why are individuals renouncing U.S. citizenship, and lately in greater numbers? I study the citizenship decision within an option value framework, arguing that for those living abroad maintaining U.S. citizenship is akin to holding an option to return to live or work in the U.S. While living abroad, U.S. citizens incur costs to maintain their U.S. citizenship, including the compliance costs associated with annual U.S. tax filing, and U.S. tax liability, if any exists after foreign tax credits and income exclusions are applied. The benefits of maintaining U.S. citizenship include the ability to return to the U.S., among other things (e.g., voting rights, traveling with a U.S. passport, the ability to pass on U.S. citizenship to one's children). The value of those benefits depends on the taxpayer's own characteristics and the characteristics of the foreign jurisdiction where they live. An example of a relevant individual characteristic is age: as individuals grow older, they may expect the probability of exercising their option to return to the U.S. to fall, and thus the value of that option to fall. An example of a relevant jurisdiction characteristic is the Rule of Law: those living in jurisdictions with more robust contract enforcement and property rights may perceive a lower probability of opting to return to the U.S. in the future, and so place a lower value on maintaining U.S. citizenship (and thus be more likely to renounce).

My theoretical option value framework motivates a regression analysis to identify the characteristics associated with the decision to renounce citizenship. I compile information on all U.S. tax filings from foreign addresses for tax years 2007-2017 and among these I identify those renouncing U.S. citizenship. I then estimate a linear probability model, and, among other results, find a significant positive effect of age on renunciation, as predicted by the option value framework. Using jurisdiction-level analysis, I find relationships similarly consistent with the option value framework: U.S. taxpayers filing from jurisdictions designated as tax havens, and with higher governance scores (measured using the World Governance Indicators' Rule of Law index), have relatively higher renunciation rates.

To address more specifically the *increase* in renunciations, I use a difference-in-difference analysis to test the effect of FATCA and related offshore enforcement efforts. These policies imposed significant costs on U.S. citizens maintaining financial accounts abroad, including compliance costs associated with new or newly enforced filing requirements, as well as potentially increased tax liabilities for those who were previously evading taxes and were induced to more accurately report foreign assets and income. Although FATCA created incentives for all Foreign Financial Institutions (FFIs) to report information on their U.S. clients, adherence to FATCA was not universal. Belnap, Thornock, and Williams (2021) show that there is variation across FFIs and jurisdictions in the quality of their information exchange with the IRS. Jurisdictions that signed Inter-Governmental Agreements (IGAs) with the U.S. settling the implementation of FATCA with their own domestic laws had higher quality information sharing. I use this variation to test the effect of compliance costs on renunciation, arguing that jurisdictions with relatively higher quality information sharing should see relatively larger increases in compliance costs for U.S. citizens living there, and thus relatively larger increases in renunciation. The data support this. IGA-signing jurisdictions had relatively larger increases in renunciation after FATCA was enacted. These results support the claim that FATCA and its associated effects on compliance costs caused an increase in citizenship renunciations by U.S. citizens living abroad. As I discuss later, most of those renouncing had no or little U.S. tax liability, after applying foreign tax credits and income exclusions, lending further support to the compliance cost explanation.

I also discuss the connection between recent expatriation tax law changes and the trends in renunciation. The data patterns suggest that some very wealthy and high-income individuals chose to leave the U.S. and renounce citizenship during the 2004-2008 period in anticipation of the introduction of a mark-to-market exit tax in 2008. The data also reveal a strong behavioral response to a reported net worth notch embedded in the expatriation tax system. Since 2004, renouncing and reporting net worth above $2 million has brought additional compliance costs and potentially additional tax liability. I show that there is a sharp drop-off in the number of individuals reporting net worth just above the threshold at $2 million and discuss several possible explanations for this pattern, including selective filing, misreporting, and real responses. Linked tax filings provide evidence of one such response: some individuals' pre-renunciation gifts and charitable contributions moved them from above to below the threshold, but this explains only a small portion of the observed pattern.

What are the policy consequences of recent renunciations? I use data on pre-renunciation tax liabilities to consider the direct revenue impacts of recent expatriations. I find that for most renunciations the revenue impact is probably negligible, because individuals had no or little U.S. tax liability in the years prior to expatriation. The distribution of liabilities is heavily skewed, however, such that a handful of individuals' renunciations have a non-trivial impact on revenues.

Considering more broadly the connection between citizenship and taxes, it is worth noting that the tax system is symmetric; those subject to the U.S. tax system consider whether to leave it, but those outside the system also consider whether to enter it. If the effects of the tax system on renunciation decisions apply similarly to the much larger group of individuals considering migration to the U.S., or naturalization (gaining U.S. citizenship) once in the U.S., the corresponding revenue impacts could be significant. That is, if the tax system discourages some individuals from moving to or naturalizing in the U.S., there is a corresponding cost in terms of lost tax revenue. I conclude by putting renunciations in a broader context, noting that even with the recent increase, renunciations still represent less than 1% of U.S. citizens living abroad, and that naturalizations are roughly 100 times more common than renunciations of U.S. citizenship (and closer to 1,000 times more likely if restricting attention to just those moving out of the U.S.).

The rest of the paper proceeds as follows. Section 1.2 sets up a conceptual framework for the costs and benefits of renunciation and briefly describes expatriation-related tax law, offshore financial enforcement, and related academic literature. Section 1.3 describes the data underlying the subsequent analyses. Section 1.4 provides a description of who is renouncing citizenship. Section 1.5 explores what can explain the recent increase in renunciations. Section 1.6 discusses the policy consequences, and Section 1.7 concludes.

1.2. Background and literature review

In this section I describe (1) the potential costs and benefits of citizenship renunciation, (2) tax law related to expatriation and how that has changed over time, and (3) tax enforcement related to offshore financial activity and how that has changed over time. Throughout the section I highlight related academic literature. For additional details on the specific steps required for citizenship renunciation, see Appendix B.

1.2.1. Costs and benefits of citizenship renunciation

The specific costs and benefits of citizenship renunciation for any given taxpayer depend on a variety of taxpayer characteristics[4], but can generally be grouped into the categories shown in Table 1.1: administrative costs and benefits (e.g., renunciation fee vs. removal of U.S. tax filing obligation) and income- or wealth-dependent tax consequences (e.g., expatriation tax consequences vs. lower future income or estate tax liabilities). This high-level framework allows a consideration of how the net benefits of renunciation would change as any of the component costs or benefits change. For example, consider one change which occurred in 2014, when the State Department raised the fee for citizenship renunciation from $450 to $2,350. This change

[4] For example, whether a taxpayer already lives or holds citizenship abroad; the amount and type of income a taxpayer receives currently and expected to receive in the future; the amount and type of assets a taxpayer holds currently and expects to bequeath in the future; the tax system of the anticipated destination country; and whether a taxpayer is currently compliant on their U.S. taxes.

uniformly lowered the net benefits of citizenship renunciation for all individuals considering it by $1,900.

Table 1.1: Costs and benefits of tax-informed citizenship renunciation

Type	Costs	Benefits
General	Administrative costs of act of expatriation (e.g., time, renunciation fee) Loss of benefits of U.S. citizenship (e.g., visa-free travel to many countries)	Reduction of ongoing administrative burden (e.g., banks wary of U.S. citizens) Reduction of yearly administrative burden (e.g., U.S. tax filing)
Income-dependent	Expatriation tax consequences	Lower future income tax liabilities
Wealth-dependent	Expatriation tax consequences	Lower future estate and gift tax liabilities

Some of these costs and benefits are simple to value (the renunciation fee is known and is exactly $2,350) while others are longer-term and more uncertain (e.g., comparing expected U.S. income tax liability vs. foreign income tax liability on the next 10 years of income). However, even when exact values are unavailable, as long as one can characterize the sign of the change, it is possible to elicit a prediction about the effect of a policy change on the incentive to expatriate. In later sections I will discuss several changes to expatriation tax law or offshore financial enforcement and consider how these policy changes would be expected to affect incentives for certain types of taxpayers considering citizenship renunciation.

1.2.2. Citizenship and U.S. tax law

The U.S. tax system has attempted to discourage tax-motivated expatriation for several decades. The Foreign Investors Tax Act of 1966 introduced §877 of the Internal Revenue Code (IRC), requiring taxation of former citizens for ten years following expatriation if tax avoidance was a "principal purpose of the expatriation" (Craig 2012). Thirty years later, a formal test for tax-motivated expatriation was introduced, as part of the Health Insurance Portability and

Accountability Act of 1996 (HIPAA). Under the new objective standards, expatriating individuals were deemed "covered expatriates" if either past-five-years average net income tax liability exceeded a certain threshold, or if net worth exceeded a different threshold.[5] Taxpayers also had to certify that they were compliant on all federal tax obligations for the five tax years preceding expatriation. As before, designation as a covered expatriate meant a taxpayer was liable for U.S. taxes on U.S.-source income and on income effectively connected with a trade or business in the U.S., at the same progressive rates faced by U.S. citizens, for the ten years following expatriation. In practice, even if a taxpayer was deemed a covered expatriate under the objective tests, one could appeal this designation and most who did so were successful.[6] Also of note, in an attempt to further discourage tax-motivated expatriation, HIPAA required the names of expatriating individuals to be published in the Federal Register (Internal Revenue Code, §6039G).

The American Jobs Creation Act (AJCA) of 2004 brought additional changes: (1) the removal of expatriates' ability to challenge their designation as tax-motivated, (2) an increase in the net worth threshold from $622K to $2M; and (3) requiring the filing of Form 8854 to complete expatriation for tax purposes.[7] The next changes were introduced in the 2008 Heroes Earnings Assistance and Relief Tax (HEART) Act, which created IRC §877A and changed the

[5] The thresholds during 2019 were $168K (average past-five-years income tax liability) and $2M (net worth). Appendix Figure 1.25 shows how these have changed over time. Note that the income tax liability threshold is applied to *tax liabilities*, not *incomes*; to have an income tax liability of $168K in 2019 would have required income of more than $500K. This distinction is sometimes missed in discussion of the expatriation tax system, with some suggesting that the threshold applies to income itself (and thus implying that many more individuals would be treated as covered expatriates according to this threshold than is truly the case).

[6] Between 1997 and July 2002, 270 applications for private letter rulings overturning the presumption of tax-motivated expatriation were made to the IRS. Of these about half received favorable responses, and all but 11 of the remainder received neutral responses. Favorable and neutral responses meant that applicants could proceed without fear of further IRS enforcement under the expatriation tax regime. This suggests that roughly 96% of appeals were successful (259/270 = 0.959) (Kwong 2009, 421).

[7] Arsenault (2009) provides further information on the first two changes. For the Form 8854 filing requirement, see the amendment history of IRC §7701(n); the 2004 AJCA added §7701(n), stating that an expatriating individual is still treated as a citizen or resident of the U.S. until that individual "provides a statement in accordance with Section 6039G."

consequences for covered expatriate designation to now include a mark-to-market exit tax, rather than the taxation of next-ten-years' U.S.-source income. Under the new regime, gains on all of a covered expatriate's assets (with a few minor exceptions[8]) are deemed realized as of the day prior to the expatriation date, and taxes owed on deemed gains above a certain exempted amount.[9] The 2008 bill also removed the requirement that Form 8854 be filed to complete expatriation for tax purposes.[10] Selected aspects and changes to the expatriation tax system are shown in Table 1.2.

Table 1.2: Selected aspects of and changes to the U.S. expatriation tax system

Expatriation date	Test for tax-motivation	Tax consequences	Other consequences
On or before June 3, 2004	Net worth > $622K (2004); Avg. inc. tax liability > $124K (2004); Presumption only, can challenge	For 10 years: taxed on U.S.-source income; estate and gifts subject to U.S. taxation	180-day limit on U.S. visits
June 4, 2004 to June 16, 2008	NW > $2M; AITL > $139K (2008); Conclusive test, cannot challenge	Same as above	Annual filings with $10K penalty for non-filing; 30-day limit on U.S. visits
On or after June 17, 2008	NW > $2M; AITL > $168K (2019)	Exit tax: mark-to-market capital gains tax (deemed realization) with $725K exemption (2019)	Annual filings until exit tax obligations are met

Notes: The column "Test for tax-motivation" indicates the tests which are applied to an individual who expatriates during the given time period; if an individual is deemed to be a "covered expatriate" under the tests, then the corresponding consequences (tax and other) apply.

Academic research on expatriation has mainly appeared in law journals, and generally focuses on detailed components of related legislation or proposed changes to the expatriation tax system (Arsenault 2009, Kwong 2009, Manolakas and Dentino 2012, Craig 2012). Westin (2000) provides a comprehensive overview of the expatriation tax system prior to the reforms of the 2000s. More recently, Ahn (2015) studies the HEART Act and notes an increase in expatriations following the introduction of the deemed realization tax that can be seen in public data from the Federal Register.

[8] Exceptions include deferred compensation items, specified tax deferred accounts, and interest in non-grantor trusts.

[9] For expatriations during 2019 the first $725K of gains are exempt. Appendix Figure 1.25 shows how the exempted amount has changed over time.

[10] Expatriating individuals are still required to file Form 8854 under IRC §6039G, but after the 2008 HEART Act's removal of IRC §7701(n), failure to file Form 8854 no longer carries the consequence that an individual is treated as a U.S. citizen or resident for tax purposes until the form is filed. This change lowered the cost of non-filing and may help explain the large share of expatriating individuals in recent years without Form 8854 filings.

Mason (2016) provides a thorough evaluation of various arguments for and against citizenship-based taxation. In response to Mason, Kim (2017) argues in favor of citizenship taxation and discusses how citizenship renunciation rates for the U.S. compare to other high-income countries. Noting the difficulty of defining a denominator when calculating the renunciation rates, Kim provides several plausible estimates based on 2010 and 2013 foreign diaspora data and relying on aggregate counts of renunciations, and concludes that the U.S. is not a serious outlier.[11] Kim also notes that "we lack empirical studies on the specific motivation of renunciation," a concern also raised by Kudrle (2015). This is precisely the gap that this paper aims to fill. More recently, De Simone, Lester, and Markle (2020) study how U.S. individuals responded to FATCA. Although their paper focuses on portfolio investments based in foreign tax havens, the authors also make use of the public Federal Register data to plot the annual counts and suggest that the recent rise in U.S. expatriations could be related to FATCA.

This paper is the first to study in detail and quantitatively the connection between citizenship renunciation and citizenship-based taxation. There is a related literature in economics which studies the connection between taxes and migration, generally studying residence-based taxation (Mirrlees 1982, Kleven, Landais and Saez 2013, Akcigit, Baslandze and Stantcheva 2016, Kleven, Landais and Muñoz, et al. 2020). The distinction between residence-based and citizenship-based taxation is important because changing one's residence is more reversible than changing one's citizenship (and may carry different costs as well). By using IRS data including Form 8854 filings, which allow for a more detailed study of the population of those renouncing citizenship, this paper

[11] Kim's estimates of renunciation rates show that the highest rates were in jurisdictions with military draft systems, with the top three rates observed for South Korea, Singapore, and Taiwan. While the relative comparison of rates across jurisdictions is certainly of interest, the many factors influencing citizenship decisions make it difficult to draw conclusions from these cross-jurisdiction comparisons. By focusing on the decisions of individuals specifically with respect to U.S. citizenship, observing trends over time, and using individual microdata, much can be learned about the motivation for citizenship renunciation and its connection to the tax system.

makes an important contribution to measuring and understanding the incentives to maintain or renounce citizenship under a citizen-based taxation system.

1.2.3. Tax enforcement and foreign financial activity

In the last decade, major changes have been made in the enforcement environment affecting financial activity by U.S. citizens living or holding financial accounts abroad. Johannesen et al. (2020) describe the introduction since 2008 of "a range of enforcement initiatives targeting owners of offshore accounts": *ad hoc* legal action and information exchanges; bilateral treaties; and FATCA.

Ad hoc legal action against Swiss banks included so-called "John Doe summonses", which allowed the IRS to request information from foreign banks about their U.S. citizen customers without identifying the specific customers in advance.[12] The IRS was authorized to use these summonses beginning in July 2008 against UBS, and subsequently against other large banks including HSBC and Credit Suisse. In addition to the *ad hoc* legal steps, the U.S. government signed bilateral information exchange agreements with several countries deemed to be tax havens.[13] These agreements allowed the IRS to request foreign bank account information for specific taxpayers in tax evasion cases. As Johannesen et al. note, citing Sheppard (2009), these agreements are relatively restrictive, requiring specification of taxpayer identities in advance and evidence to justify the request, and thus may not be effective deterrents of offshore tax evasion.

Finally, a new reporting regime requiring systematic information exchange on U.S. citizen account holders between foreign financial institutions (FFIs) or foreign tax authorities and the IRS

[12] If required to specify customers in advance, the IRS would not have been able to meaningfully pursue the relevant information. U.S. taxpayers hiding assets did not notify the IRS of their holdings, and thus could not be identified *ex ante* and specified in requests for information.

[13] Between 2008 and 2010, the U.S. signed such agreements with six jurisdictions: Liechtenstein, Luxembourg, Malta, Monaco, Panama, and Switzerland.

was introduced in 2010, as part of FATCA. This may have affected U.S. citizens living abroad in two main ways. First, the IRS would now have better access to third-party reporting on income and assets for these individuals. Second, these individuals now faced increased costs (either financial costs or compliance costs) in their dealings with FFIs, as those FFIs themselves faced increased costs in complying with FATCA. Dharmapala (2016) studies how a unilateral reporting regime (like FATCA) affects the cost to FFIs of providing financial services and how this in turn affects incentives for tax-compliant behavior by foreign residents. Belnap, Thornock, and Williams (2021) study foreign countries' and FFIs' participation in automatic information sharing with the IRS and show that FFI participation was near-universal (97% of FFIs participated in automatic information sharing) and costly. They also show, however, that the quality of information shared varies considerably across FFIs and jurisdictions.

These enforcement changes are relevant to the study of citizenship renunciation because each change either made it more difficult, or less attractive, to be a U.S. citizen living and maintaining financial accounts abroad. This paper is the first to study carefully the potentially unintended consequence of these changes in tax enforcement – increased U.S. citizenship renunciation by U.S. citizens living abroad.

1.3. Data

The main source of data for this study is an IRS database of former U.S. citizens who have renounced their citizenship since 1998. Individuals who expatriate are required to meet with a consular official, resulting in a Certificate of Loss of Nationality (CLN), and to file a form with the IRS (Form 8854, the Initial and Annual Expatriation Statement, intended to be filed along with the income tax filing for the year of expatriation). The State Department notifies the IRS of each CLN, which the IRS then matches with the Form 8854 filings they receive from taxpayers. In practice, some individuals have only one of the two forms, and the IRS database represents the union of renouncing individuals based on CLNs, Form 8854s, or both. In this paper I study only those renunciations occurring between 1998 and 2018, to allow for a lag in 8854 filing and ensure a more complete picture of the renunciations occurring in each year. The database also includes information about some of the individuals relinquishing long-term residency status (rather than U.S. citizenship). Because this information is not entirely complete—not all such individuals are included in the database—I restrict my focus in this paper to former citizens.

For all individuals in the database, I observe the date of renunciation and the destination country or jurisdiction. For individuals with Form 8854 filings I observe reported net worth as of the date of expatriation. Other fields of interest on Form 8854 that are not available for study at this time include more details on how foreign citizenship was acquired, as well as a breakdown of assets by asset category. For those with Social Security Numbers (SSN) or Taxpayer Identification Numbers (TIN), I can link to their other relevant tax filings. About 70% of those renouncing have these identifiers. I include all individuals in each analysis where possible, although at times this is not feasible (e.g., when studying pre- renunciation income, which requires linking to income tax filings).

In addition to the database of information about those renouncing citizenship, I also have access to the complete database of all U.S. tax filings. I use data from several forms in particular, including Form 1040 (Income Tax), Form 1116 (Foreign Tax Credit), Form 2555 (Foreign Earned Income Exclusion), and Form 709 (Gift Taxes). I rely especially on information about the population of U.S. tax filers who are filing from abroad.[14] This allows me to observe the base of individuals residing abroad who could potentially renounce their U.S. citizenship.

[14] I am especially grateful to Tom Hertz at the IRS for developing these data.

1.4. Description of renouncers

This section answers the first of my three research questions: Who is renouncing? Previous studies have had to rely exclusively on publicly available information, which in practice has meant only the names of individuals expatriating each quarter as reported in the Federal Register. I provide more detailed information on these individuals, including their prior U.S. tax filing behavior (and the resulting inferred location, i.e., in the U.S. or abroad), self-reported net worth and income, and destination jurisdictions.

1.4.1. Overall counts

Figure 1.1 shows the annual count of all former citizens who have renounced citizenship, as identified in the IRS database, from 1998 to 2018. There is a gradual increase in annual counts during the 2000s, followed by a more marked increase since 2011. This is the pattern of renunciations that was available for study prior to this paper, using only publicly available information about those expatriating.[15]

[15] Appendix Figure 1.26 shows the annual count using publicly available information, with counts for 1962-1994 from the Joint Committee on Taxation (1995) and counts for 1998-2020 from the Federal Register.

Figure 1.1: Annual count of U.S. citizenship renunciations

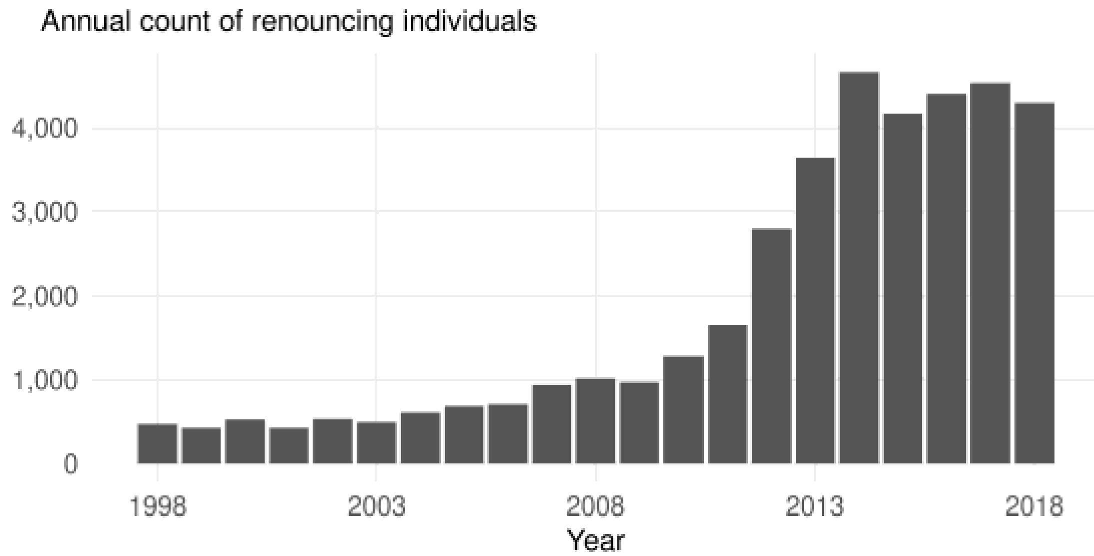

Notes: This figure plots the count of former citizens who renounced citizenship each year, as identified in the IRS database for years 1998-2018.

1.4.2. Prior presence in the U.S.

Presented with the overall increase, a policy-relevant question is, are these individuals "leaving the U.S.", or instead individuals who already were living abroad and chose to drop U.S. citizenship? To answer this I link individuals to their pre-renunciation income tax filings and infer their locations from the addresses reported on those filings. Most individuals are required to file Form 1040 each year, even those living abroad. I categorize each individual into one of a few buckets: those that filed at least once from a U.S. address before renouncing ("Movers"); those that filed income tax returns but never from a U.S. address ("Droppers"); and those for whom we cannot observe pre-renunciation locations (either because they have no filings or have no TIN).[16] Because this method relies on data for tax filings available in the years prior to renunciation, I limit

[16] This is an imperfect proxy that in general would bias towards classification as a Mover, as some individuals may maintain addresses in the U.S. even while living abroad, or may use a U.S.-based tax preparer's address on their filings. Note that because not all renouncing individuals are primary filers, I search for tax filings associated with their TIN as either primary or secondary filers, to ensure I gather as much pre-renunciation location information about each individual as possible.

this classification to those renouncing in 2005 or later[17]; I then use five years of pre-renunciation tax returns to classify each individual as Movers or Droppers.

Figure 1.2 shows the count of renouncing individuals each year, split by this classification. The gray bars represent the Movers – those who can be thought of as "leaving the U.S.". The orange bars represent the Droppers, those who were filing returns but always from a foreign address. In blue are those with a TIN but no filings, or without TINs or SSNs to match to tax returns (this latter group is likely comprised mainly of Droppers, i.e., those who were not present in the U.S. prior to expatriation, which would explain why they have no filings or no TINs). While the annual count of Movers has increased slightly, most of the of the recent increase is by Droppers. In later sections I will study further what can explain this increase in Droppers, arguing that it is primarily an unintended consequence of the increased compliance costs resulting from FATCA and other offshore financial enforcement.

[17] The IRS database of income tax returns starts in earnest with returns for tax year 1998.

Figure 1.2: Annual count of renunciations, split by pre-renunciation tax filing locations

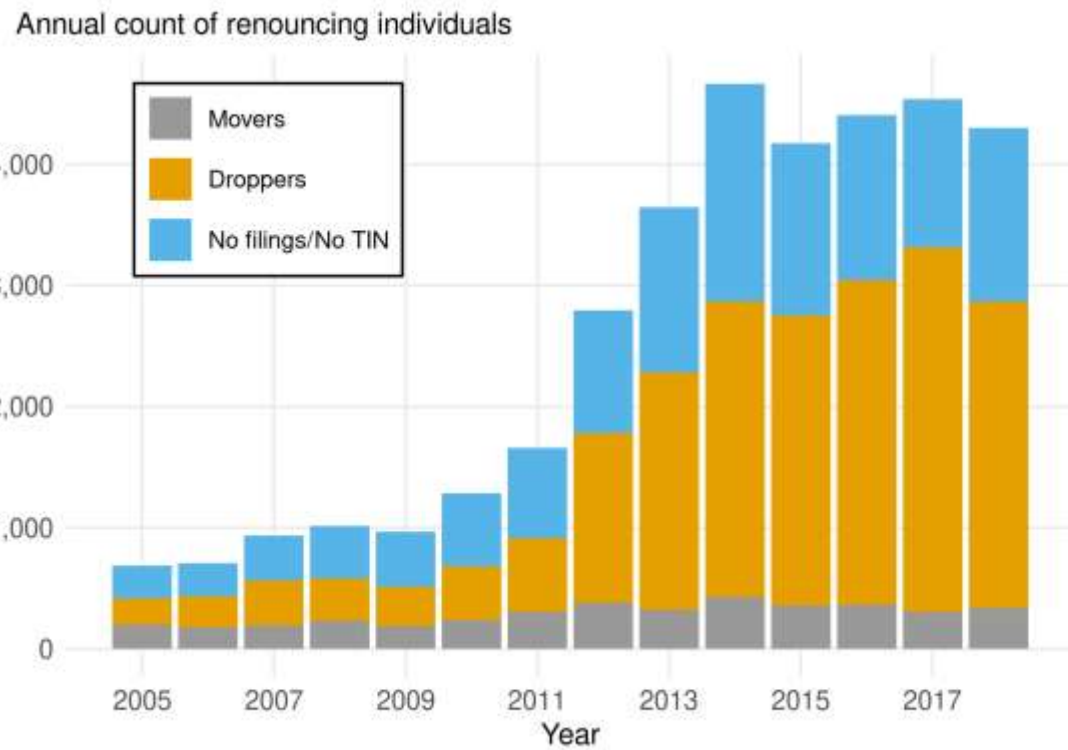

Notes: This figure plots the count of individuals renouncing each year, split by their classification based on Form 1040 filing behavior in the five years prior to renunciation. Movers are those who filed at least once from the U.S. during those five years; Droppers are those who filed always from abroad. Renunciations prior to 2005 are excluded to ensure sufficient pre-renunciation data are available.

1.4.3. Income and wealth

It is also interesting to consider how these individuals compare to others in terms of income and wealth. I begin by comparing renouncers to other foreign filers and the full population of U.S. tax filers, in terms of total and wage income, and then compare income within renouncers, between Movers and Droppers. I then do a similar comparison for reported net worth.

Figure 1.3 reports the mean values of total income and wage income in the year prior to renunciation, and compares this to two other groups: (i) all other filings from foreign addresses, and (ii) a sample of the full population of Form 1040 filings. In orange are renouncers who were the primary filer for a linked 1040 in the year prior to renunciation.[18] In blue are all other Form 1040 filings from foreign addresses for the given year, and in gray are a sample of all Form 1040 filings. The vertical dashed lines represent three key dates related to expatriation tax law: 2004 AJCA (raising the net worth threshold for covered expatriate designation), 2008 HEART Act (introducing the mark-to-market exit tax), and 2010 FATCA (increasing information reporting of foreign financial accounts held by U.S. citizens). To illustrate the influence of a few outliers on the mean value among renouncers, the dashed line removes the top 10 individuals for each year.

Figure 1.3: Comparison of income for renouncers, foreign filers, and all tax filers

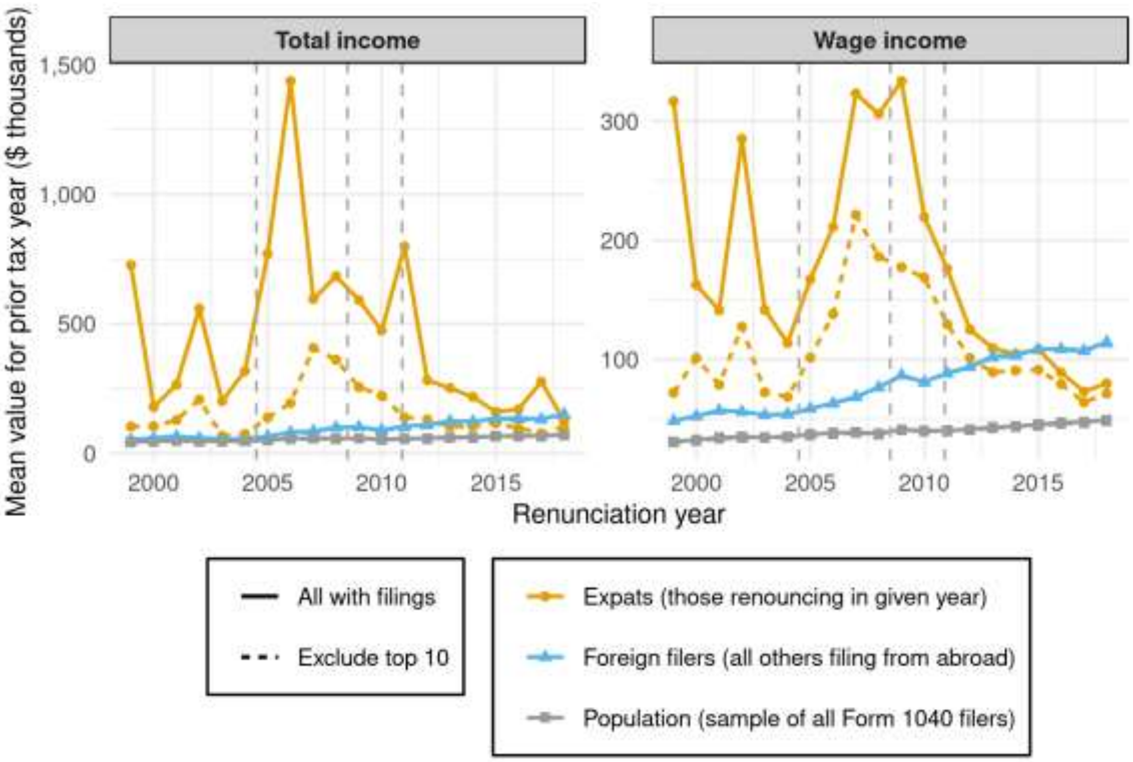

Notes: This figure compares the income of renouncers in the year prior to renunciation to two comparison groups: all other foreign filings, and a sample of the population of Form 1040 filing. For renouncers, only primary filers with linked filings are included. The three vertical dashed lines represent three key dates related to expatriation tax law: the 2004 AJCA, the 2008 HEART Act, and 2010 FATCA. The solid line includes all individuals; the dashed line removes the top 10 in each year.

[18] I use the prior year to ensure a full year's income is reported. In the year of renunciation itself, those renouncing citizenship file a Form 1040 representing the portion of the year they are a citizen, and may file a Form 1040 NR for the remaining portion of the year after they have renounced.

Figure 1.4 demonstrates that the average income of those renouncing each year has changed dramatically over time, and that outlier individuals play an important role in driving the annual averages. Those renouncing during the window between 2004 (AJCA) and 2010 (FATCA) were on average much higher income, relative to those expatriating in the 2010s; and this is true even when removing the top 10 individuals each year. Prior to 2010, those renouncing were higher income, on average, than both other foreign filers and the broader U.S. filer population. Since 2010, those renouncing have had lower income, on average, than other foreign filers, but still higher than the U.S. filer population overall. Similar trends appear when considering the median values instead of the mean (see Appendix Figure 1.15).

For more detail about the income of renouncers, consider Figure 1.4, which compares the income just for renouncers, with averages calculated separately for Movers and Droppers.[19] For both groups, incomes were higher during the 2005-2010 time period, but the big outliers for total income are among the Movers, not the Droppers. The groups also differ in terms of their source of income; Movers have higher average total income, but Droppers have higher average wage income. The dramatic influence of the top 10 individuals each year on the average total income among Movers is a stark example of the nature of the renunciation policy problem: although most individuals have a small revenue impact, a handful can have a significant effect; I discuss this in further detail in Section 1.6. As above, similar trends are seen in the median values (see Appendix Figure 1.16).

[19] Those without filings or TINs are excluded due to lack of income data.

Figure 1.4: Comparison of income among renouncers, Movers vs. Droppers

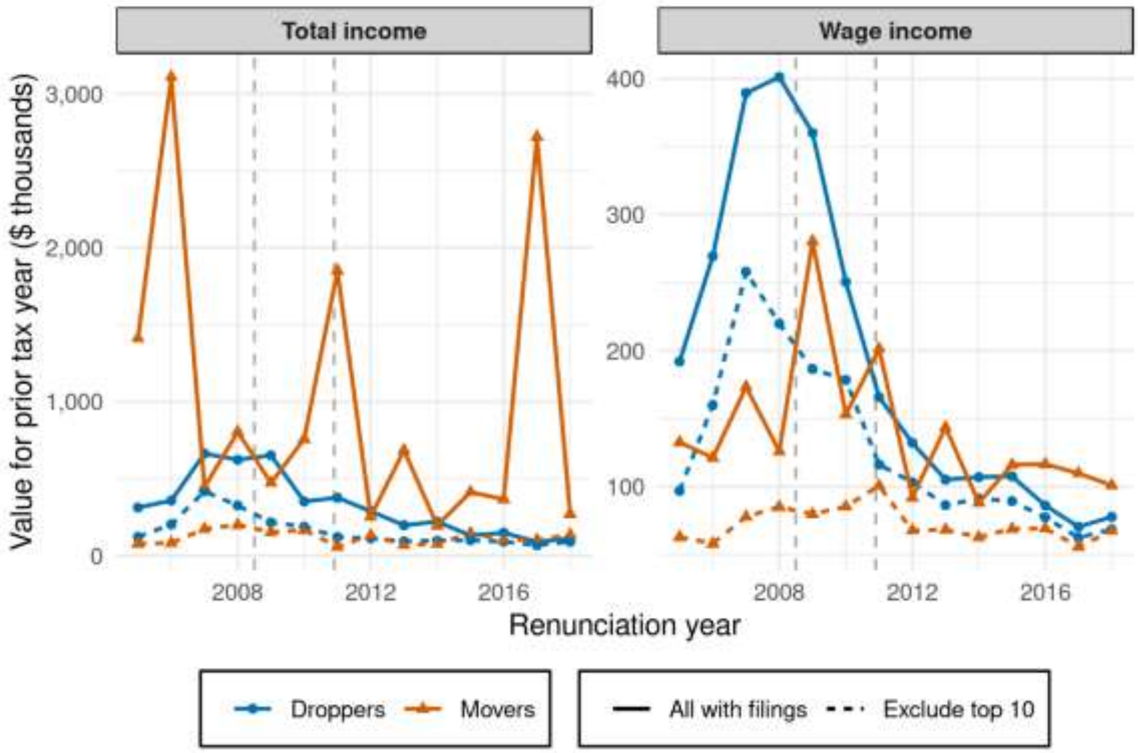

Notes: This figure compares the income in the year prior to renunciation for renouncers with linked Form 1040 filings as primary filers. The mean values are calculated separately among Movers and Droppers. Renouncers with no filings or no TINs are excluded.

Moving from income to wealth, I begin by grouping the renouncers based on their net worth as reported on Form 8854. I construct buckets using the thresholds for covered expatriate designation: $622K (the threshold prior to the AJCA, i.e., prior to June 2004) and $2M (the threshold since the AJCA, i.e., after June 2004). Figure 1.5 shows the annual count, grouped by reported net worth.[20]

[20] Reported net worth is only completely available since mid-2004, when Form 8854 began to require all filers to list their reported net worth; prior to this change, only those with net worth above the tax-motivation threshold ($622K in early 2004, adjusted upward for inflation over 1998-2004) were required to report this information.

Figure 1.5: Annual count of renunciations, split by reported net worth

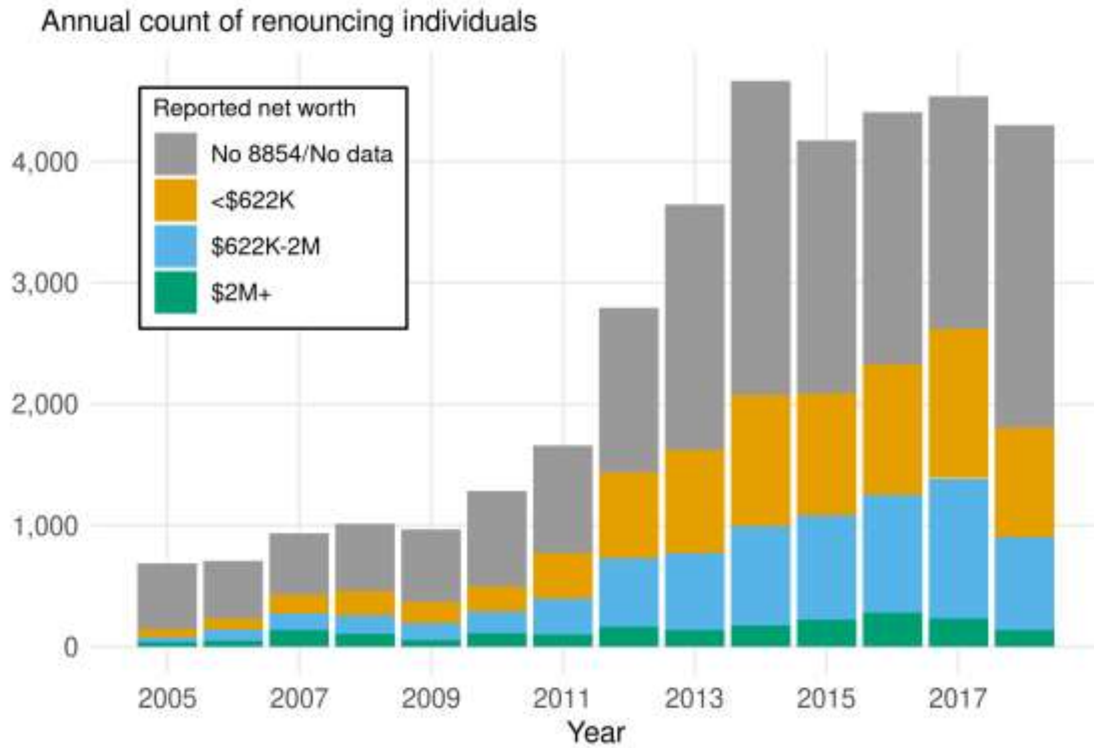

Notes: This figure plots the count of individuals renouncing each year, split by reported net worth. Prior to June 2004, reported net worth data are not available. This figure starts with the first full year of available data, 2005.

A few patterns are worth noting. First, although there has been a small rise in the number of renunciations by those reporting net worth of at least $2M (the green bars), these still represent a relatively small share of the total. Second, there has been more substantial growth in the number reporting between $622K and $2M in net worth (the blue bars); this group is relevant because it represents the individuals who prior to the AJCA would have been designated as covered expatriates, but after the raising of the net worth threshold no longer faced such designation. At the same time, there was similar growth in those reporting less than $622K (the orange bars). Finally, an important pattern is the persistent large share of renunciations without Form 8854 or without reported net worth data, (the gray bars). Although filing Form 8854 is a necessary step to fully complete one's citizenship renunciation, a significant number of individuals still have not done so. Some of this pattern in more recent years could reflect that some file Form 8854 with a lag (this likely explains the difference between 2017 and 2018 – those who renounced in 2018 and plan to file Form 8854 may still be finalizing their filings). Although this non-filing limits the ability to draw comprehensive conclusions about the wealth of all renouncers, useful information can still be gleaned by studying those for whom data are available.

I next consider how the wealth of those renouncing each year has changed over time, and whether this differs for Movers and Droppers. Figure 1.6 reports the mean reported net worth of those renouncing each year since 2005, separately for Movers and Droppers (only including those with reported net worth data available). The patterns are similar to those above for income: Movers are wealthier than Droppers; average renouncer wealth during the 2004-2010 period was notably higher than in more recent years; and removing the top 10 individuals in each group each year has a dramatic effect on the average values. Similar patterns emerge when considering the median values (see Appendix Figure 1.17).

Figure 1.6: Comparison of reported net worth among renouncers, Movers vs. Droppers

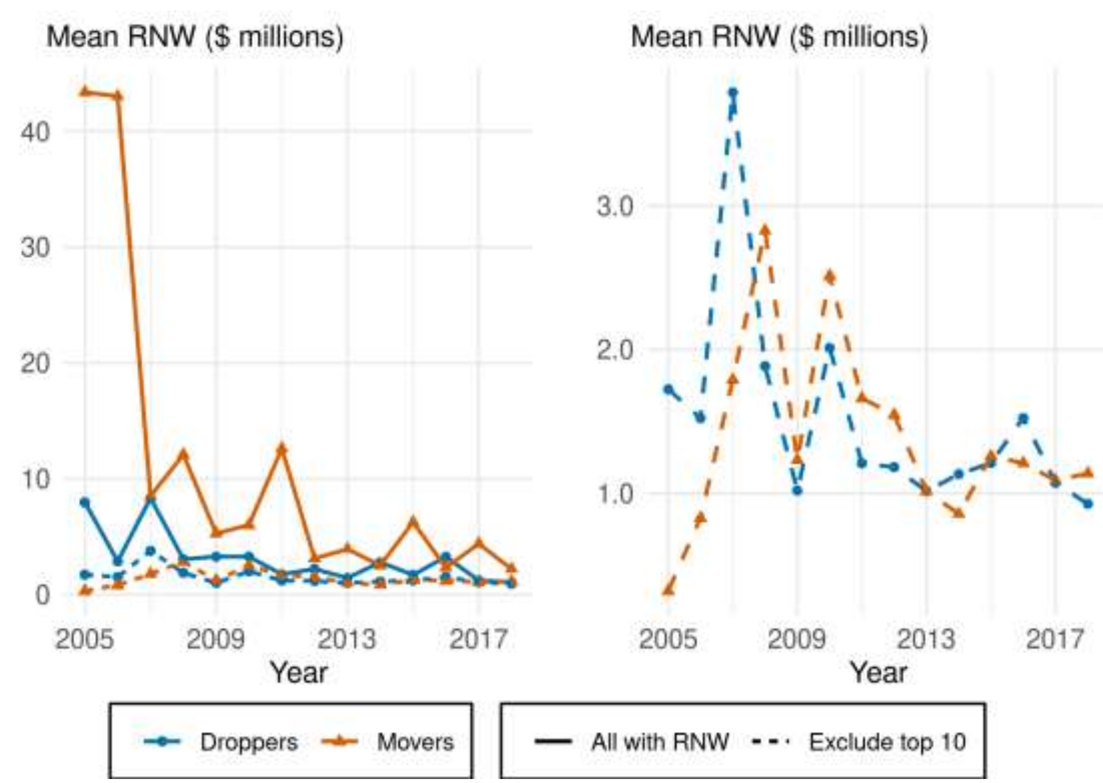

Notes: This figure compares reported net worth among those renouncing each year, separately for Movers and Droppers. Only those with reported net worth data available are included. The left panel includes all Movers and Droppers; right panel drops the top 10 Movers and Droppers, by reported net worth, each year.

Finally, I consider how the wealth distribution among renouncers compares to the population. Table 1.3 shows the count of renouncers from 2005-2018 by their reported net worth, as well as the total reported net worth in each group. The share of the population in each net worth group is included, based on the 2019 Survey of Consumer Finances (these population estimates are for households, and thus weighted towards higher amounts, relative to the renunciation statistics which are for individuals). I provide two estimates for the share of renouncers in each net worth group: the first assumes that all renouncers without reported net worth data are in the <$1M group; the other excludes those without reported net worth data (i.e., it assumes those without reported net worth data are distributed the same as those with data).

Table 1.3: Comparison of reported net worth groups

	Expatriates			Population	Expatriates		Expatriates
Reported net worth	Number	Share, assuming missing are <$1M	Share, excluding missing	Share (housholds)	Total reported net worth ($B)	Share of total	Median age
<$1M	10,700	82.7%	63.4%	88.1%	$3.90	8.0%	47
$1-2M	4,240	11.8%	25.1%	5.7%	$6.18	12.7%	56
$2-10M	1,430	4.0%	8.5%	5.1%	$6.20	12.8%	53
$10-100M	470	1.3%	2.8%	1.0%	$13.76	28.3%	51
$100M+	50	0.1%	0.3%	0.1%	$18.52	38.1%	45
Has 8854, no RNW	1,860						47
No 8854, no RNW	17,040						45
Total	35,790	100.0%	100.0%	100.0%	$48.56	100.0%	

Notes: This table reports statistics for individuals who renounced between 2005 and 2018. Renouncer counts are rounded to the nearest 10 for disclosure purposes. Population share is based on household shares in the 2019 Survey of Consumer Finances.

Renouncers are relatively wealthier than the population. Specifically, millionaires are relatively more common: 17% of renouncers (assuming none of those missing data are millionaires) versus the estimate of 12% among households in the U.S. population; and note that estimates for the U.S. population share of millionaires among *individuals* are lower, around 5-

10%.[21] While those renouncing are on average wealthier than the population, the numbers also reveal the relatively small scale of ultra-wealthy expatriations. Between 2005 and 2018 only about 50 renouncers reported net worth above $100 million. However, although small in number, these individuals may have an outsize impact on policy; their decisions to expatriate tend to show up in the news and spur legislative changes.[22] Interestingly, above $1 million, the median age at expatriation decreases with reported net worth.

Taken together, the information on income and wealth shows that those who have chosen to renounce citizenship were on average higher income and higher wealth than the population, but this average obscures significant heterogeneity within the renouncer population: a few outliers in each year strongly influence the average values. The pattern over time shows that average income and wealth among renouncers has been trending down.

1.4.4. Destination jurisdictions

Finally, I provide information about renouncers' destination jurisdictions. "Destination" is perhaps a misnomer given that many of these individuals always lived in the foreign jurisdiction or moved there many years prior to dropping U.S. citizenship. Nonetheless, destination here refers to the foreign jurisdiction listed as an individual's country of tax residency (when reported) or general residency (when tax residency is not reported or available).[23] Renouncers' destinations are

---

[21] The 2018 Credit Suisse Global Wealth Report estimates that 17.35 million Americans were millionaires, or 7.1% of the adult population.

[22] For example, legislative changes in the 1990s reportedly came about because then President Bill Clinton read about the tax-motivated expatriation of six wealthy Americans in Forbes magazine **Invalid source specified.**. More recently, Senators Chuck Schumer and Bob Casey proposed a bill to punish Facebook co-founder Eduardo Saverin for his pre-Facebook IPO expatriation **Invalid source specified.**. The bill, titled the Expatriation Prevention by Abolishing Tax-Related Incentives for Offshore Tenancy, or Ex-PATRIOT Act, failed to make it out of committee.

[23] In almost all cases, tax residency and general residency are the same: more than 99% of the records with both tax residency and general residency have the same jurisdiction reported for both.

of interest generally, and may also provide some information about whether taxes are an important factor in the expatriation decision.

Figure 1.7 shows the share (in Panel A) and count (in Panel B) of renouncers in each year going to the top five destination jurisdictions (by total count from 1998-2018), and all others. Over time renunciation has become more concentrated in the top five destinations, with the share going to destinations outside these top five falling from about 50% in the 2000s to 30% in 2013, although this share ticked back up to 40% by 2018. In recent years, the share renouncing to Canada has risen dramatically. Also of note is the sharp rise and gentler fall in renunciations to Switzerland.

Figure 1.7: Share and count of renunciations to top destination jurisdictions

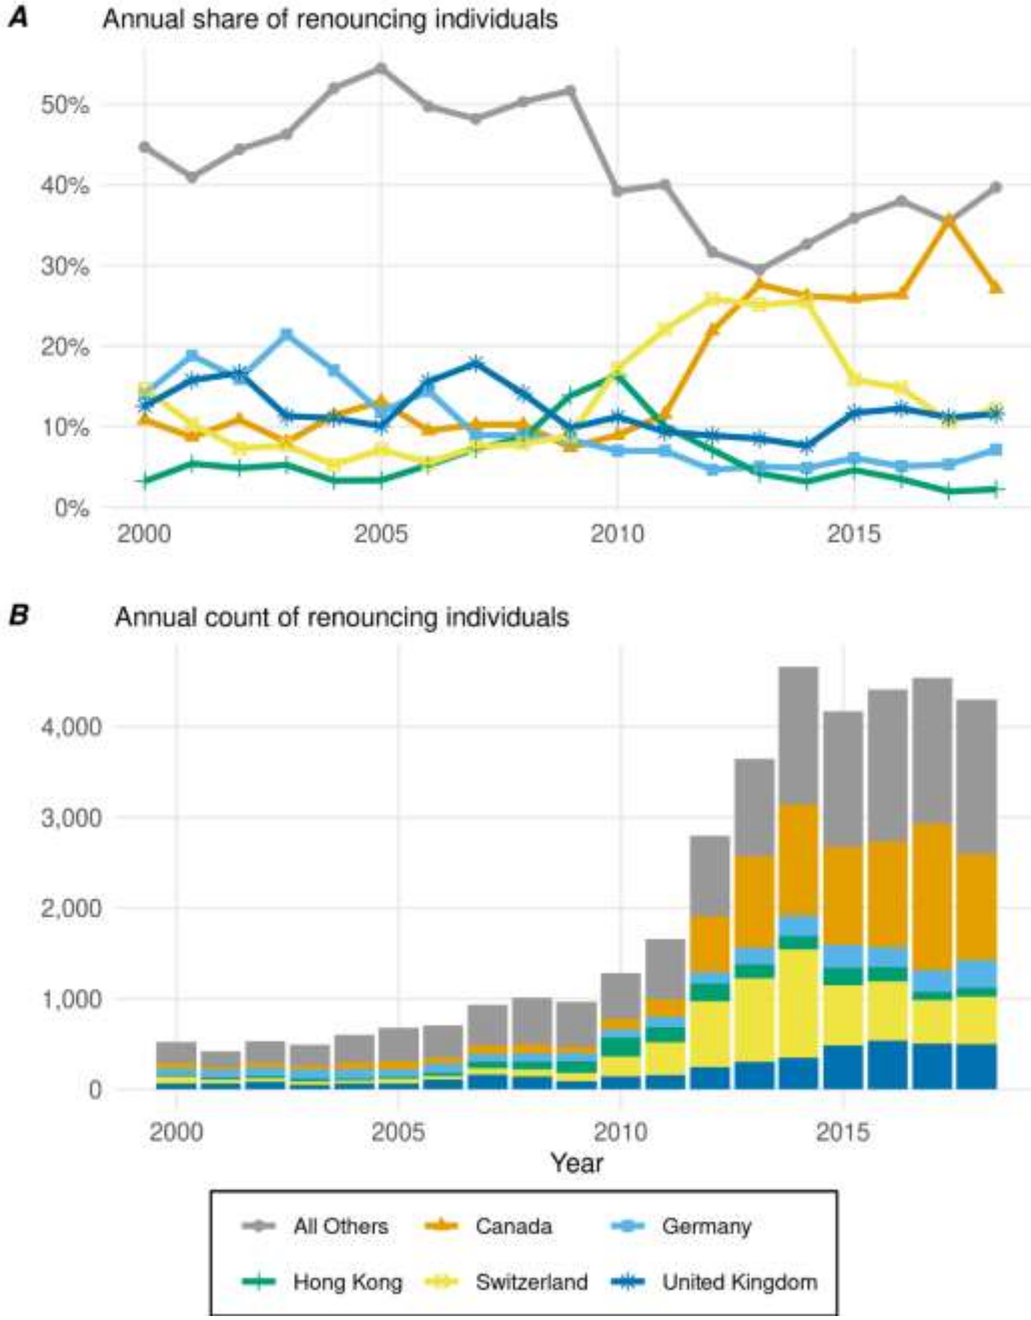

Notes: Panel A plots the share of individuals in each year renouncing to each of the top five jurisdictions, or all others. Panel B plots the count of individuals renouncing to each of these jurisdictions, or all others.

I next consider how the pattern of renunciations to certain jurisdictions relates to the base of U.S. citizens filing from those jurisdictions. If renunciation were equally likely regardless of where a U.S. citizen living abroad is located, then the number of U.S. citizens filing from a jurisdiction should correlate perfectly with the number of U.S. citizens dropping their citizenship in that jurisdiction. To test whether the data follow such a pattern, I construct two rankings: first, I rank foreign jurisdictions by the average number of U.S. tax filings received each year from each jurisdiction; second, I rank the same foreign jurisdictions by the average number of renunciations each year listing that jurisdiction as their destination. I then produce a scatterplot of these rankings. I do this exercise separately for the years 2007-2010, and 2011-2018, in order to test whether the patterns change before and after FATCA.[24]

Figure 1.8 shows the rank-rank plots described above. A few patterns are worth noting. First, most jurisdictions fall close to the 45-degree line, suggesting that the correlation between foreign filings and renunciation is strong, on average. Second, there are clusters of jurisdictions that fall away from the 45-degree line. Above the line are jurisdictions whose renunciation rank is higher than their foreign filer rank; U.S. citizens filing from these jurisdictions are more likely to renounce citizenship, on average, than those filing from other jurisdictions. The prevalence of tax havens among these clusters suggests that tax considerations do play a role in some citizenship decisions.[25] Those below the line are jurisdictions where renunciation is less common than would be expected, based solely on the number of foreign filings. The difference between the pre- and post-FATCA plots also suggests the composition of renouncers may have changed between the two time periods. I study these patterns further in Section 1.5.1.

[24] At present I have comprehensive data on foreign filings by year and jurisdiction only since tax year 2007. If in future these data are available for earlier years, one could extend this analysis to include those additional years.

[25] In this discussion, and later in Section 1.5.1, I rely on the list of tax havens used in Johannesen et al. (2020). As they note in footnote 1, "This list does not have any official role in IRS enforcement efforts; the IRS does not have an officially accepted definition of a tax haven."

Figure 1.8: Comparing jurisdictions by renouncer rank vs. foreign filer rank

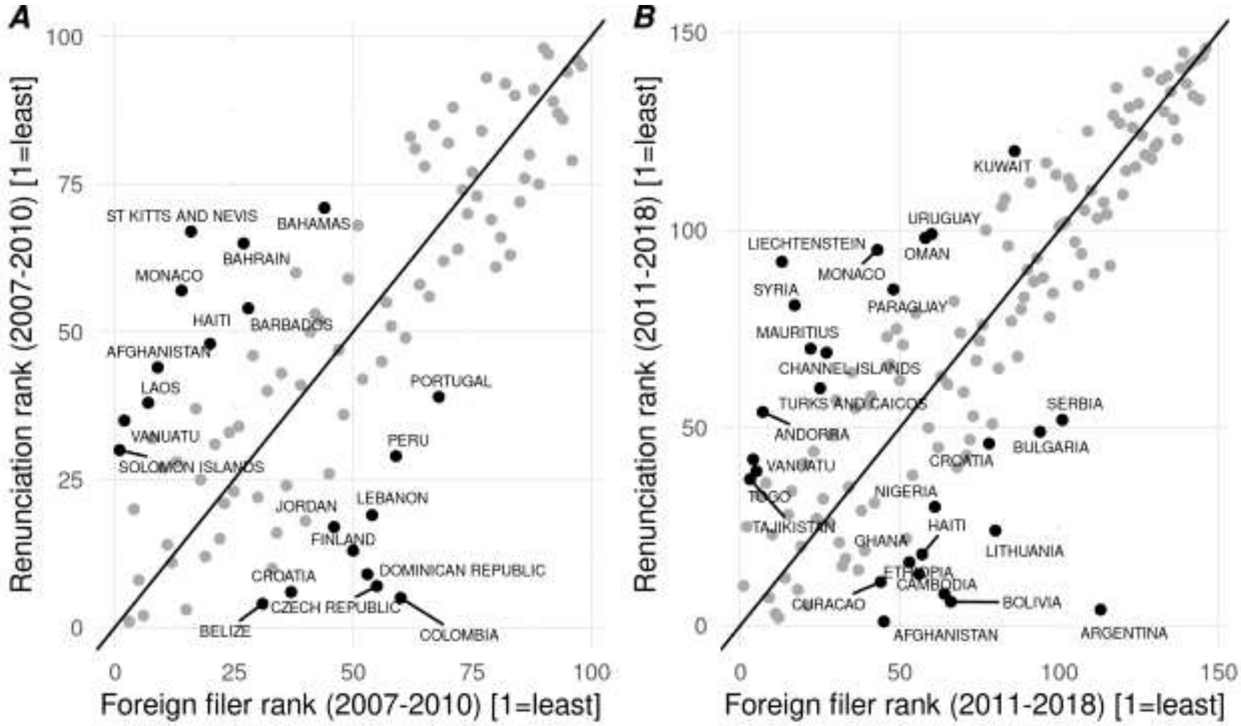

Notes: This figure plots each jurisdiction's rank based on renunciations (the average annual count of individuals reporting the jurisdiction as their destination when renouncing) and foreign filings (the average annual count of U.S. tax filings received from the jurisdiction). The ranks are calculated separately for the years 2007-2010 and 2011-2018 to test whether patterns change before and after FATCA.

One might ask whether the top destinations differ when focusing on particular sub-groups (e.g., the wealthy). In general, this is not the case; although there are some small differences, the top jurisdictions are consistent when looking within various subgroups. In Appendix A, Table 1.7 shows the top ten destinations within each reported net worth group; Table 1.8 the top ten destinations within each renouncer classification (Mover vs. Dropper); and Table 1.9 shows the top ten jurisdictions ranked by total renunciation counts, and by renunciations as a share of unique U.S. tax filers. Tax havens are more common in the ranking by share.

1.5. Explaining the increase in renunciations

This section addresses the second of my three research questions: Why are individuals renouncing? I focus first on explaining the recent increase, which as shown above is mainly driven by Droppers. I then consider renunciations by Movers and their connection to U.S. expatriation tax policy.

1.5.1. Explaining renunciations by those already living abroad

I develop a framework for the decision of those living abroad to maintain or drop citizenship using a simple option value approach. I then use this framework to motivate empirical tests, first using individual-level data to test various determinants of renunciation and confirm that age is positively correlated with renunciation, then using jurisdiction-level data to test what characteristics correlate with greater renunciation frequency, and finally using a difference-in-differences approach to show that increased compliance costs help explain the recent increase in renunciations.

1.5.1.1. Theoretical framework

For U.S. citizens living abroad, U.S. citizenship can be thought of in an option value framework. For those abroad, U.S. citizenship represents an American-style call option in which the foreign resident U.S. citizen retains the right to return to the U.S. to live or work at some point in the future. Typically, option value can be decomposed into time value and intrinsic value. Time value for the option on U.S. citizenship corresponds to age: as individuals get older, the remaining time in which they can exercise the option decreases, leading the value of that option to decrease as well. All else equal, this suggests that the probability of renunciation should increase with age.

The intrinsic value of the option on U.S. citizenship comprises many components. First, consider that for a typical financial option, the value of that option increases with the volatility of the underlying asset. Similarly, the value of U.S. citizenship should increase as volatility increases. Volatility in this case could include global economic uncertainty and the political stability of foreign countries relative to the United States; those living in more stable countries may consider themselves less likely to want or need to exercise the option to return to or work in the U.S., and thus be more likely to renounce U.S. citizenship. Other components of the intrinsic value could include the tax rates of the foreign country relative to the U.S. and the relative value of the foreign country's passport. For those living in countries with lower relative rates, the value of the option on U.S. citizenship would be lower, while for those in countries with a relatively more valuable passport, the option value of being able to use one's U.S. passport would be lower.

Finally, in addition to the value of the option, consider the cost of maintaining it. This cost has always included remitting one's annual tax liability, if any, as well as the compliance costs, including time and effort, of annual filing of U.S. tax returns. These compliance costs have increased in recent years, with additional forms required for many taxpayers, both by tax agencies and financial institutions. In the next sections I test whether the predictions of this framework are borne out in the data.

1.5.1.2. Individual determinants of renunciation

To begin testing the implications of the options model, I focus first on identifying characteristics associated with the costs and benefits of the decision of those living abroad to renounce citizenship. Although not all the reasons someone might choose to renounce are captured in tax filings, administrative microdata still allow me to test how several key characteristics relate to renunciation.

The base for this study is the set of all U.S. tax filings by those filing from abroad. This includes Form 1040 filings, and other linked tax form data, for those filing from abroad for tax years 2007-2017. Among these filings, I identify the individuals who ultimately renounce citizenship, and flag the tax year prior to the year in which they expatriate, dropping subsequent filings for these individuals if they appear.[26] As noted above, I consider the tax information in the year prior to the year of expatriation as the most relevant, because it is represents a complete year of earnings and other taxpayer decisions. The final dataset contains about 17,000 instances of citizenship renunciation (I include only primary filers, and am unable to include individuals without TINs or linked tax filings), out of more than four million tax filings from those living abroad.

I develop a simple linear probability model, regressing *Renounce* (the decision to renounce citizenship in the following year) on a set of individual-year covariates and, in some specifications, jurisdiction, year, or jurisdiction X year fixed effects:

$$Renounce_{ijt+1} = \beta(Covariates_{it}) + [\alpha_j] + [\alpha_t] + [\alpha_{jt}] + \varepsilon_{ijt}$$

These covariates include: total positive income (TPI) in millions of dollars; wages as a share of total positive income (0 if no TPI); a dummy indicating the taxpayer had a positive tax liability (net of foreign tax credits and income exclusions); dummies indicating whether a taxpayer had nonzero values reported for Schedule C or Schedule E income, respectively[27], a dummy indicating that a charitable contribution deduction was claimed on Schedule A, a dummy indicating Form 709, the U.S. Gift and Generation-Skipping Transfer Tax Return, was filed; and a dummy indicating a taxpayer received any notice from the IRS. In some specifications I include age (in years), though this slightly lowers the observation count because of some missing data on dates of birth. Appendix Table 1.10 presents summary statistics for these variables.

[26] Some individuals who expatriate continue to file Form 1040 or Form 1040 NR after renunciation, depending on their income sources and other circumstances.

[27] Schedule C includes income and loss from a business or profession practiced as a sole proprietor; Schedule E includes income and loss from rental real estate, royalties, partnerships, S corporations, estates, trusts, and residual interest in real estate mortgage investment conduits (REMICs).

The results of the basic linear probability model are shown in Table 1.4.[28] The dependent variable is coded as 100 or 0, so that the coefficient estimates represent the effect in percentage points for each covariate, holding all others constant. The different columns include various combinations of fixed effects, culminating in column (6) with year X jurisdiction fixed effects included (so that the model seeks to explain the decision to renounce within a jurisdiction in a year). Figure 1.18 in the Appendix plots the coefficient estimates, scaled by the mean probability of renunciation, to show the estimated percent change in the probability of renunciation resulting from a 0 to 1 change in each binary covariate. The figure also compares the coefficient estimates when including or excluding Movers, showing similar coefficient estimates.

[28] In this main specification, seeking to explain the recent increase in Droppers, I include only the Droppers as renouncers, excluding Movers from the dataset in any year where they appear. I also run the models including all renouncers, and the results are nearly identical; see Appendix A, Table 1.11.

Table 1.4: Individual linear probability model results

Dependent variable:	Binary: Renounce in following year (100/0)					
	[1]	[2]	[3]	[4]	[5]	[6]
Total Positive Income	0.0102	0.0102	0.0108	0.0069	0.0074***	0.0071
($ millions)	(0.0064)	(0.0063)	(0.0064)	(0.0056)	(0.0018)	(0.0055)
Wage share	-0.0844	-0.055	-0.1119	-0.1203*	-0.1086***	-0.1044**
(% of TPI)	(0.0783)	(0.0603)	(0.0736)	(0.0549)	(0.0075)	(0.0386)
Positive tax liability	-0.0805	-0.0834	-0.0695	-0.0724	-0.0662***	-0.0645
(1/0)	(0.0641)	(0.0650)	(0.0630)	(0.0466)	(0.0059)	(0.0446)
Any Sch C income	0.0886***	0.0937***	0.0672**	0.0597***	0.0482***	0.0518**
(1/0)	(0.0248)	(0.0278)	(0.0229)	(0.0177)	(0.0088)	(0.0176)
Any Sch E income	0.0066	0.0042	-0.0104	0.0135	-0.0052	0.0033
(1/0)	(0.0372)	(0.0383)	(0.0366)	(0.0269)	(0.0081)	(0.0262)
Schedule A charity	-0.0136	-0.0218	0.0176	-0.1208	-0.0984***	-0.0963
(1/0)	(0.0881)	(0.0856)	(0.0903)	(0.0926)	(0.0112)	(0.0926)
Filed gift tax return	2.2867***	2.2807***	2.2498***	2.2302***	2.1869***	2.1767***
(1/0)	(0.4204)	(0.4172)	(0.4143)	(0.4277)	(0.0551)	(0.4175)
Received any notice	0.1070**	0.1023*	0.0313	0.1163*	0.0398***	0.0447
(1/0)	(0.0544)	(0.0547)	(0.0197)	(0.0529)	(0.0080)	(0.0280)
Age		0.0019			0.0022***	0.0022**
(years)		(0.0013)			(0.0002)	(0.0009)
Constant	0.3954***	0.2890***				
	(0.1356)	(0.0712)				
Year FE	No	No	Yes	No	Yes	No
Jurisdiction FE	No	No	No	Yes	Yes	No
YearXJurisdiction FE	No	No	No	No	No	Yes
Observations	4,831,000	4,790,000	4,831,000	4,831,000	4,790,000	4,790,000
Adjusted R^2	0.0005	0.0005	0.0013	0.0045	0.0053	0.0066
Mean dep. var.	0.343	0.343	0.343	0.343	0.343	0.343

Notes: *p<0.1; **p<0.05; ***p<0.01. Standard errors, clustered by year and by jurisdiction, are shown in parentheses. The dependent variable is coded as 100 or 0 so that the coefficient estimates represent the effect in percentage points for each covariate, holding all others constant. "Movers" that can be linked to Form 1040 filings as a primary filer are excluded here; results when including them are shown in Appendix Table 1.11.

The results suggest several individual characteristics connected with the decision to renounce citizenship. Filing a gift tax form, which is relatively rare in general, is very strongly associated with renunciation (consistent with a pattern I demonstrate later related to the net worth threshold for covered expatriate designation). The presence of Schedule C income is positively associated with renunciation, while having a higher wage share of income is negatively associated with renunciation (interesting given the pattern shown in Section 1.4.3 that Droppers had relatively high wage income, suggesting that those renouncing had both high wage income and non-wage income). Having a positive tax liability is *negatively* associated with expatriation, although this is only statistically significant at standard levels in one specification. However, if the association is truly negative, this would be consistent with an explanation in which long-term foreign resident U.S. citizens drop citizenship because of increased compliance costs (filing new and more complicated forms), not because of tax liability itself.

Most relevant to the option value framework, the results show that age is significantly, and positively, correlated with the decision to renounce. This is consistent with the prediction that as individuals age, the time value of their option on U.S. citizenship decreases, leading to lower values for that option and renunciation becoming more common.

1.5.1.3. Jurisdiction characteristics and renunciation frequency

The previous section tested whether certain individual characteristics, observable in tax filings, correlate with the decision to renounce in a way consistent with the option value framework. In this section, I similarly test whether characteristics of the *jurisdictions* from which foreign-resident U.S. citizens file their taxes correlate with the prevalence of renunciations from those jurisdictions. The option value framework predicts a higher value of U.S. citizenship (and thus a lower rate of renunciation) for those living in foreign jurisdictions where they perceive a higher probability of wanting or needing to exercise the option by returning to live or work in the United States.

For this test, I collapse the individual-level data to a jurisdiction-level dataset and estimate the following equation:[29]

$$Renunciation\ share_j = \beta\left(Covariates_j\right) + \varepsilon_j$$

The renunciation share is defined as the total number of renunciations in a given jurisdiction from 2008 to 2018, divided by the unique set of U.S. tax filers from that jurisdiction over the period 2007 to 2017. The denominator approximates the set of "potential renouncers" – those who filed from abroad and could have chosen to renounce U.S. citizenship. Dividing the total number of renunciations to a jurisdiction by this set of potential renouncers gives an outcome value that allows comparison of the relative frequency of renunciation across jurisdictions.

The covariates are motivated by the option value framework. First are three binary variables indicating whether a jurisdiction is designated as a tax haven, relying on the designations in Johannesen et al. (2020); offers citizenship-for-sale, based on Christians (2017); and is majority native English-speaking. Next, I include separately the average percentile rank of the jurisdiction on two measures from the World Governance Indicators: the Rule of Law and Political Stability indices (higher values indicate better governance). I also include the average percentile rank of each jurisdiction's passport value according to the Henley Passport Index, a ranking of passports

[29] I take this approach to focus specifically on the jurisdiction characteristics and to capture associations over a longer time period (collapsing across years), relative to the individual approach above. However, I also test the relationship between the jurisdiction characteristics and the probability of renunciation by merging the characteristics into the individual-level data and running similar specifications to those in the prior section, replacing the jurisdiction fixed effects with the characteristics I discuss in this section. The results are generally consistent between the two approaches, though the two are not directly comparable: the individual approach studies the decision to renounce in a given year, while the jurisdiction approach studies the frequency of renunciations over a longer time period. See Appendix Table 1.13 for the results of the individual-level regression with jurisdiction-level covariates.

based on the number of destinations accessible without a prior visa (higher values indicate a more valuable foreign passport). Finally, I include the average annual change in real GDP, according to the IMF. Summary statistics are shown in the Appendix Table 1.12.

The results are shown in Table 1.5 below. Tax haven jurisdictions are associated with higher renunciation shares, consistent with lower taxes motivating renunciation for at least some individuals. The Rule of Law index is also positively correlated with renunciations, consistent with the option value framework's prediction that individuals living in more stable jurisdictions anticipate a lower likelihood of exercising the U.S. citizenship option, and thus are more likely to renounce. Similarly, jurisdictions with more valuable passports are associated with higher renunciation shares. In these specifications, citizenship-for-sale (CFS) is negatively correlated with renunciation share; I also run specifications excluding CFS, or including a tax haven X CFS interaction (see Appendix Table 1.14). Removing CFS does not materially affect the other covariate estimates, and the haven interaction results suggest that the CFS effect is driven by the few non-haven CFS jurisdictions, like Bulgaria and Serbia, where renunciation is relatively uncommon. The lack of an effect for the Political Stability index reflects the strong positive correlation between the Rule of Law and Political Stability indices. Overall, the results are generally supportive of the predictions of the option value framework.

Table 1.5: Jurisdiction-level regression results

	Dep. var.: Total renunciations/unique foreign filers			
	[1]	[2]	[3]	[4]
Tax haven (1/0)	0.0262** (0.0125)	0.0266** (0.0129)	0.0353** (0.0156)	0.0171* (0.0088)
Citizenship-for-sale (1/0)	-0.0095 (0.0064)	-0.0112* (0.0065)	-0.0155* (0.0082)	-0.0068 (0.0053)
English-speaking (1/0	-0.0056 (0.0073)	-0.0121 (0.0087)	-0.0157 (0.0097)	-0.0052 (0.0059)
Rule of Law index (percentile)		0.0323*** (0.0093)	0.0182** (0.0072)	0.0201*** (0.0062)
Political Stability index (percentile)		-0.0032 (0.0069)	0.0015 (0.0078)	-0.0042 (0.0051)
Passport ranking (percentile)			0.0120** (0.0057)	0.0088 (0.0055)
Change in Real GDP (percentage points)				0.0002 (0.0005)
Constant	0.0091*** (0.0010)	-0.0038 (0.0033)	-0.0047 (0.0036)	-0.0033 (0.0030)
Observations	213	205	196	187
Adjusted R^2	0.0852	0.1707	0.2134	0.2253

Notes: *p<0.1; **p<0.05; ***p<0.01. Heteroskedasticity-robust standard errors are shown in parentheses.

1.5.1.4. Testing the compliance cost explanation

I next seek to explain the *increase* in renunciations seen in the past decade. Many public press articles about the recent increase include anecdotes attributing the increase to FATCA and its associated compliance costs for U.S. citizens living abroad; academic articles have posited this explanation as well (e.g., Kudrle (2015), De Simone, Lester and Markle (2020)). These articles note that over the past decade there has been a general increase in offshore financial enforcement, including FATCA, as well as *ad hoc* legal and information actions and bilateral treaties that, for certain countries, increased the flow of information to the IRS about U.S. citizens' financial assets and earnings abroad. In general this meant that foreign financial institutions (FFIs) faced increasing compliance costs when working with U.S. citizen customers, and they became less willing to do so.[30] Thus for those already living abroad, maintaining U.S. citizenship in the 2010s brought additional costs, including difficulty dealing with local financial institutions and more strictly enforced filing requirements.[31] If one wished to remain abroad, these costs were only avoidable by dropping U.S. citizenship.

To test empirically to what extent increased compliance costs explain the recent increase in renunciations, I exploit variation across jurisdictions in the implementation of FATCA-related information sharing. Although FATCA provides strong incentives for all FFIs to report on their U.S. citizen clients, adherence to FATCA varies across FFIs and across jurisdictions, as demonstrated by Belnap, Thornock, and Williams (2021) (henceforth "BTW"). The authors show that jurisdictions that signed Inter-Governmental Agreements (IGAs) with the U.S. settling the implementation of FATCA with their own domestic laws had higher quality information sharing.[32] BTW also show that information sharing is costly for FFIs, and relatively more costly for those

[30] Press reports describe numerous anecdotes of U.S. citizens abroad facing such difficulties. See, e.g., Williams **Invalid source specified.**, "U.S. expats find their money is no longer welcome at the bank" and Graffy **Invalid source specified.**, "The law that makes U.S. expats toxic." Some of these difficulties are only now starting to arise, as FATCA implementation was not necessarily immediate; France, for example, was set to start reporting information in 2020, prompting an August 2019 article warning of pending bank account closures for 40,000 U.S. citizens **Invalid source specified.**.

[31] One important group of individuals who were particularly affected by the enforcement changes were those hiding assets abroad. These individuals faced an ever-increasing likelihood of being discovered by the IRS. One response to this would be to come clean, pay any necessary penalties, and maintain U.S. citizenship. Another response would be to drop U.S. citizenship in an attempt to "sneak out" before the hidden assets could be discovered. However, because hidden assets are unobservable it is not possible to test directly whether individuals with such assets were more likely to expatriate following the increased enforcement actions.

[32] There are two model types of IGAs. Under Model 1, firms report to their local tax authorities, who then report to the IRS. Under Model 2, firms report information from their consenting account holders' accounts directly to the IRS, while information from non-consenting account holders is sent through their local tax authorities. The majority of signed IGAs follow Model 1. BTW show specifically that information exchange quality is higher for Model 1 jurisdictions relative to Model 2 and Non-IGA jurisdictions. For these tests, I consider both all IGAs together, and the two Model types separately.

that report higher quality information.[33] Therefore, I hypothesize that U.S. citizens living in IGA-signing jurisdictions, where information exchange is likely of higher quality, should experience larger compliance cost impacts from FATCA, and thus be relatively more likely to renounce after FATCA.

Figure 1.9 below shows graphically that after FATCA, renunciation rates increased in IGA-signing jurisdictions, but stayed flat in non-IGA jurisdictions. Furthermore, the increase is driven by Model 1 jurisdictions, with renunciation rates staying flat for Model 2 jurisdictions (though at a higher level relative to non-IGA jurisdictions). The figure plots, for each year, the number of renunciations as a share of prior-year tax filers from each group. The years 2008-2010 are the "pre" period to the left of the red line indicating the passing of FATCA in December 2010, while 2011-2018 are the "post" period. In this figure, and in the following regression analysis, I exclude observations from Switzerland because of the unique, pre-FATCA focus on tax avoidance and evasion by U.S. citizens there, although this does not affect the conclusions.

[33] Dharmapala (2016) models the behavior of FFIs under FATCA, with FFIs passing on the costs of information sharing to their accountholders through increased fees.

Figure 1.9: Annual share renouncing for IGA and non-IGA jurisdictions

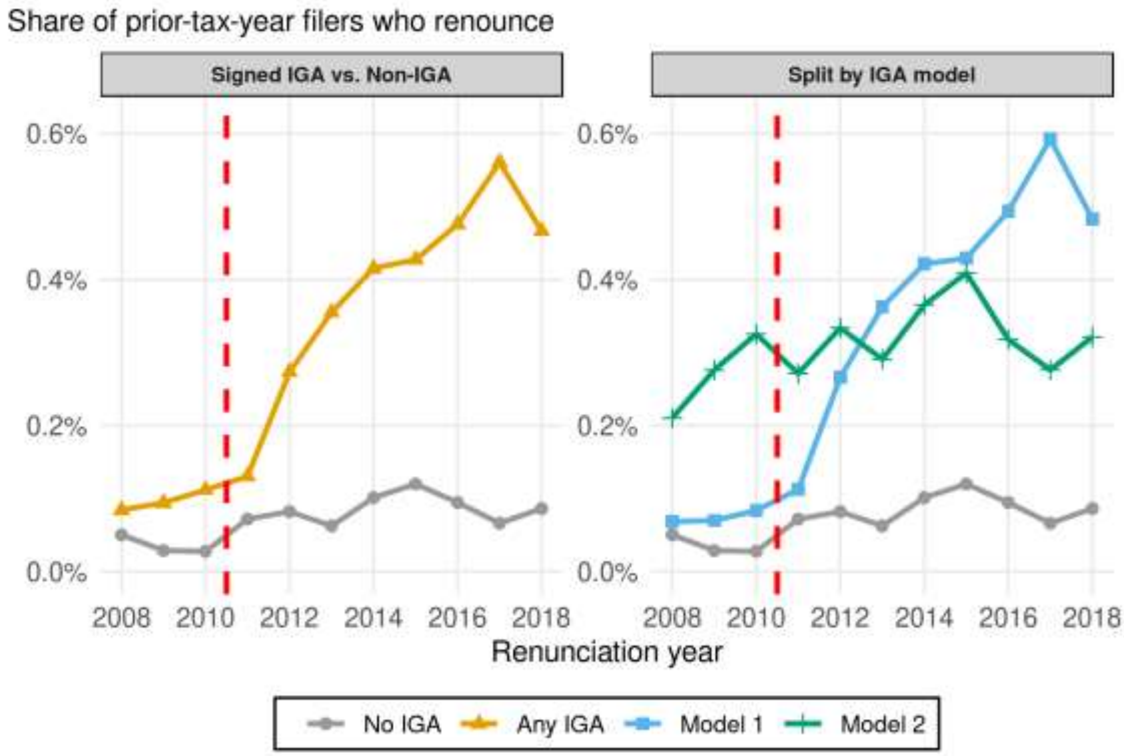

Notes: This figure plots the share of individuals renouncing (share of prior-tax-year filers), based on the data underlying the individual regressions and splitting the sample based on IGA designations.

I also test this relationship using a difference-in-differences approach, relying on the same individual-year level data as in the previous section, as follows:

$$Renounce_{ijt+1} = \alpha + \beta_1 IGA_j * Post_t + \beta_2 IGA_j + \beta_3 Post_t + \gamma(Covariates_{it}) + \varepsilon_{ijt}$$

where IGA is an indicator equal to one for jurisdictions that signed an IGA, $Post$ is an indicator for tax years 2010 or later (renunciations in 2011 or later), and $IGA * Post$ is their interaction. I also run this model including separate indicators and interactions for Model 1 and Model 2 IGAs. As before, *Renounce* is an indicator equal to 100 if an individual renounces citizenship the following year, and 0 otherwise.

The main regression results are shown in Table 1.6 below and confirm that the probability of renunciation increased relatively more after FATCA for individuals living in IGA-signing jurisdictions. This result is driven by Model 1 jurisdictions; Model 2 jurisdictions saw no such post-FATCA increase.

Table 1.6: Difference-in-difference results

Dep. var.: Renounce in following year (100/0)	[1]	[2]
Post	0.0455*	0.0454*
	(0.0251)	(0.0251)
IGA jurisdiction	0.0465	
	(0.0338)	
Post X IGA	**0.2579****	
	(0.1011)	
Model 1		0.0217
		(0.0240)
Post X Model 1		**0.2899*****
		(0.1050)
Model 2		0.2363
		(0.1686)
Post X Model 2		**0.0065**
		(0.0275)
Individual covariates	Yes	Yes
Observations	4,690,000	4,690,000
Adjusted R^2	0.0013	0.0014
Mean dep. var.	0.312	0.312

Notes: *p<0.1; **p<0.05; ***p<0.01. Standard errors clustered by year and by jurisdiction are shown in parentheses. IGA jurisdiction is an indicator for jurisdictions that signed a FATCA IGA in or before 2019; Model 1 and Model 2 are indicators identifying whether the model type of the signed IGA. Post is an indicator for tax years 2010 and later.

The results shown in Figure 1.9 and Table 1.6 are robust to various alternative data filters and variable definitions. Including Switzerland, excluding "Movers", or excluding jurisdictions with fewer than 100 U.S. tax filers does not affect the results (see Figure 1.21 and Table 1.15 in Appendix A). Alternative IGA definitions also do not affect the results, whether restricting to only

IGAs signed through 2015 or 2017, or also including unsigned but agreed-in-substance IGAs (see Table 1.16 in Appendix A).

Taken together, these tests provide empirical evidence consistent with the compliance cost explanation for the recent increase in renunciations. The option value framework predicts that if compliance costs increase, renunciations should increase, and the tests here show that renunciations became relatively more common after FATCA in IGA-signing jurisdictions, where resident individuals were more likely to experience increased compliance costs.

1.5.2. Evaluating the effects of expatriation tax law changes

The previous section focused on explaining the recent increase in renunciations by Droppers. What about Movers? Although in fact a small share of the total, those who at one point filed from the U.S. and subsequently moved abroad and renounced citizenship represent more of the stereotypical individual that may come to mind when thinking about expatriation and citizenship renunciation. Indeed, these are the types of individuals cited by legislators when discussing the expatriation tax system, and who in the past have apparently prompted changes to that system. Understanding their behavior is important for evaluating the effects of prior tax law changes, and considering future policy.

1.5.2.1. General trends in renunciation by Movers

As noted above, Movers may have a large set of reasons for moving abroad and renouncing their citizenship, including family or other ties abroad, but it is also possible that tax considerations play an important role in their decisions. To better understand the relationship between renunciations by Movers and the tax system, I begin by showing their annual counts; this is the solid gray line in Figure 1.10 below (the same as the gray bar in Figure 1.2 above). I then adjust this count in two ways. First, note that my preferred categorization of individuals as Movers or Droppers relies on five years of pre-renunciation tax filings, to allow for the fact that some of those moving from the U.S. take several years to settle in before renouncing. One drawback of this approach is that it limits observations to those renouncing in 2005 or later. To address this, I produce an alternate categorization based only on tax filings in the two years prior to renunciation, shown with the blue line. Second, because much of the public press and legislative focus on this subject has centered on wealthy or high-income Movers, I produce a set of counts restricting to those Movers who have high net worth (above $622K) or high income (AGI greater than $200K in the year prior to renunciation), shown with the dashed lines.

Figure 1.10: Annual count of renunciations by Movers

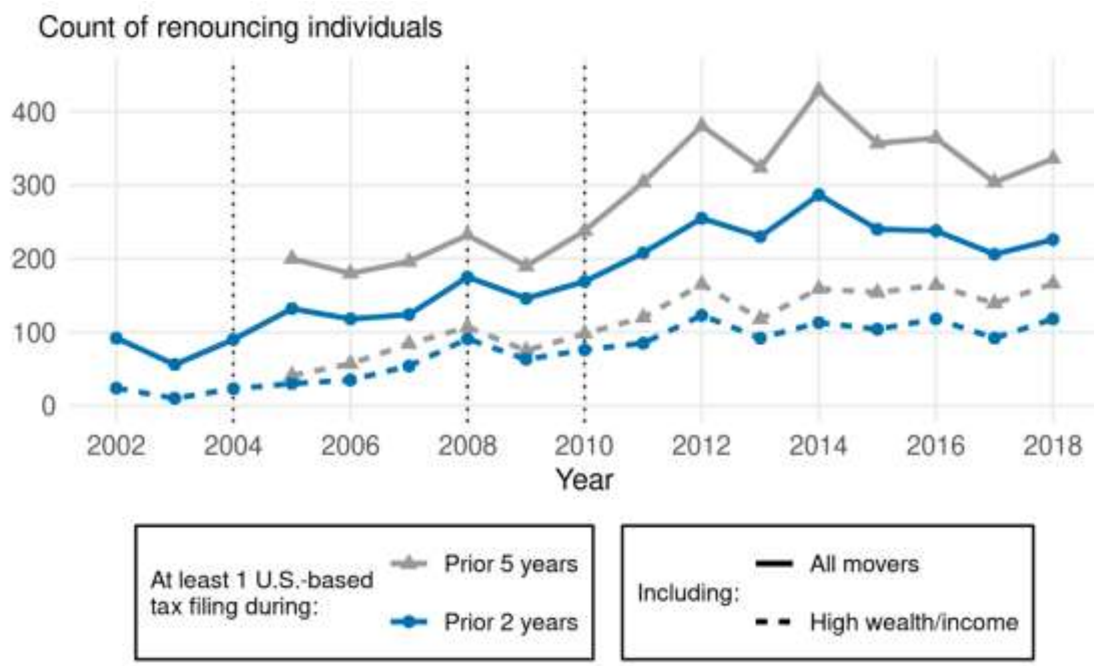

Notes: This figure plots the annual count of renunciations made by those designated as Movers, either based on having filed from the U.S. at least once during the prior five years (in gray) or two years (in blue). Total counts are shown with solid lines, and those including only high wealth (net worth > $622K) or income (prior-year AGI > $200K) are shown with dashed lines. Vertical dashed lines indicate years with legislative changes: 2004 (AJCA), 2008 (HEART Act), and 2010 (FATCA).

Over time there has been an increase in renunciations by Movers, although the increase in the past decade is less extreme than that seen above for Droppers. The number is still small, with annual counts of around 100-200 during the 2000s, and 300-400 in the 2010s. The increase in renunciations by Movers between 2004 and 2008 could indicate that the tax law changes in those years had some effect; I explore this further below. The acceleration in renunciations after 2010 suggests that the increase in offshore financial enforcement may also have played a role in the renunciation decisions for Movers, just as was seen above for Droppers; perhaps some U.S. citizens who previously would have moved abroad but maintained citizenship chose instead to renounce that citizenship when facing increased compliance costs during the 2010s.

1.5.2.2. Relating tax law changes to renunciations

How did the 2004 and 2008 tax law changes affect individuals' decisions of whether and when to renounce? The 2004 AJCA made two important changes to the expatriation tax system: (1) it raised the net worth threshold for designation as a covered expatriate from $622K to $2M; and (2) it removed the ability to challenge one's designation as tax-motivated, replacing it with a strictly objective test based on net worth, past-five-years average tax liability, and certification of compliance with the last five years of tax filings.

Consider how these two changes would affect the costs and benefits of renunciation. The net worth threshold change would lower the cost for certain individuals. For individuals with true net worth between $622K and $2M, renunciation prior to the change would have included designation as a covered expatriate and the ensuing effort to either challenge that designation or deal with the next-10-years tax consequences. After the change, these individuals could renounce and report their true net worth without being designated as covered expatriates. The change may also have affected some individuals with true net worth above $2M, as the cost of getting under the threshold was lowered.[34] Individuals with true net worth below $622K would be unaffected, as both before and after the change they were not at risk of covered expatriate designation. In sum, the change lowered the cost of expatriation for those with net worth above $622K, and especially for those above $622K and below $2M. All else equal, this predicts more renunciation by such individuals as a result of the increase in the net worth threshold.

The removal of the ability to challenge covered expatriate designation should work in the opposite direction, *raising* the cost of renunciation for individuals who previously could have successfully challenged their covered expatriate designation. Consider two wealthy individuals (above the $2M threshold), identical in every respect except that one has no ties abroad, while the other does have strong ties in the country to which they plan to renounce. Prior to this change the former individual would be designated as a covered expatriate and may have some difficulty challenging that designation; the latter would also be designated as covered but would have an easier time challenging that designation. After the 2004 removal of the ability to challenge, both individuals would be designated as covered and remain so. The effect of this removal is thus a change in the relative cost of renunciation: for the individual with strong ties abroad, the relative cost of renunciation has increased when compared to the cost for an individual without strong ties abroad. All else equal, this predicts relatively fewer renunciations by those with ties abroad, and thus relatively more by those without such ties.

The 2008 HEART Act's introduction of the mark-to-market exit tax was a more fundamental change to the expatriation tax system. It changed the consequences of covered expatriate designation from an uncertain future liability based on an income over the next 10 years with an

[34] For example, consider someone with $2.1M in true net worth; prior to the change, they would need to somehow lower their reported net worth by nearly $1.5M to fall below the $622K threshold, but only by $100K to fall below the new $2M threshold.

immediate, up-front tax liability based on unrealized capital gains above an exemption threshold (although this liability could be temporarily deferred). This change could in principle push in different directions. On the one hand, an up-front liability could be perceived as more costly than the uncertain future liability, and thus make renunciation seem more costly than under the prior system. On the other hand, the ability to pay the one-time exit tax and cleanly walk away may have been more desirable to some individuals, relative to the lingering connection to the U.S. that would persist under the next-10-years system. Whether the mark-to-market tax would be more or less desirable than the earlier system would also depend crucially on the extent of an individual's unrealized capital gains; someone with significant wealth but relatively low amounts of unrealized capital gains would face little or no liability under the mark-to-market tax, which exempts the first several hundred thousand dollars of gains. In sum, the change from the earlier system to the mark-to-market exit tax was certainly a significant change, but its effects would likely not push unambiguously in the same direction for all individuals.

1.5.2.3. Trends around the tax law changes

To further understand how the AJCA and HEART Act affected the number of renunciations, and the types of individuals renouncing, we can look for evidence in the patterns of renunciation around the tax law changes. Consider Figure 1.11, which shows the count of renouncing individuals, grouped by reported net worth, in each half-year time period from 2003-2010. Beginning with the effect of the AJCA, and focusing on the net worth threshold change, we would ideally compare the number with net worth between $622K and $2M, before and after the threshold change, to see whether under the post-AJCA regime in which they are no longer designated as covered expatriates, their numbers rise. Unfortunately, data on net worth is not available for the pre-AJCA renunciations, and even post-AJCA, many individuals either do not have a filed Form 8854 or do not have reported net worth data available.

Nevertheless, we can observe that the number of renunciations increased after the AJCA change; if the expatriate provisions in the AJCA were intended to discourage renunciations, a simple assessment of the trend suggests they may not have achieved that goal. We can also observe that after the net worth threshold was increased, there were a handful of individuals with reported net worth in the $622K to $2M range (the green bars); it is possible they were induced to renounce by no longer facing the cost of covered expatriate designation. At the same time, however, there were a similar number of renunciations by those with reported net worth above $2M (the top, light blue bars), confirming that renunciation was still desirable for some individuals even when facing the costs of covered expatriate designation. Without further detail on the net worth of all renouncers, both before and after the 2004 law change, it is difficult to draw firm conclusions about the effect of the change.

Figure 1.11: Annual count of renunciations, before and after AJCA and HEART Act changes

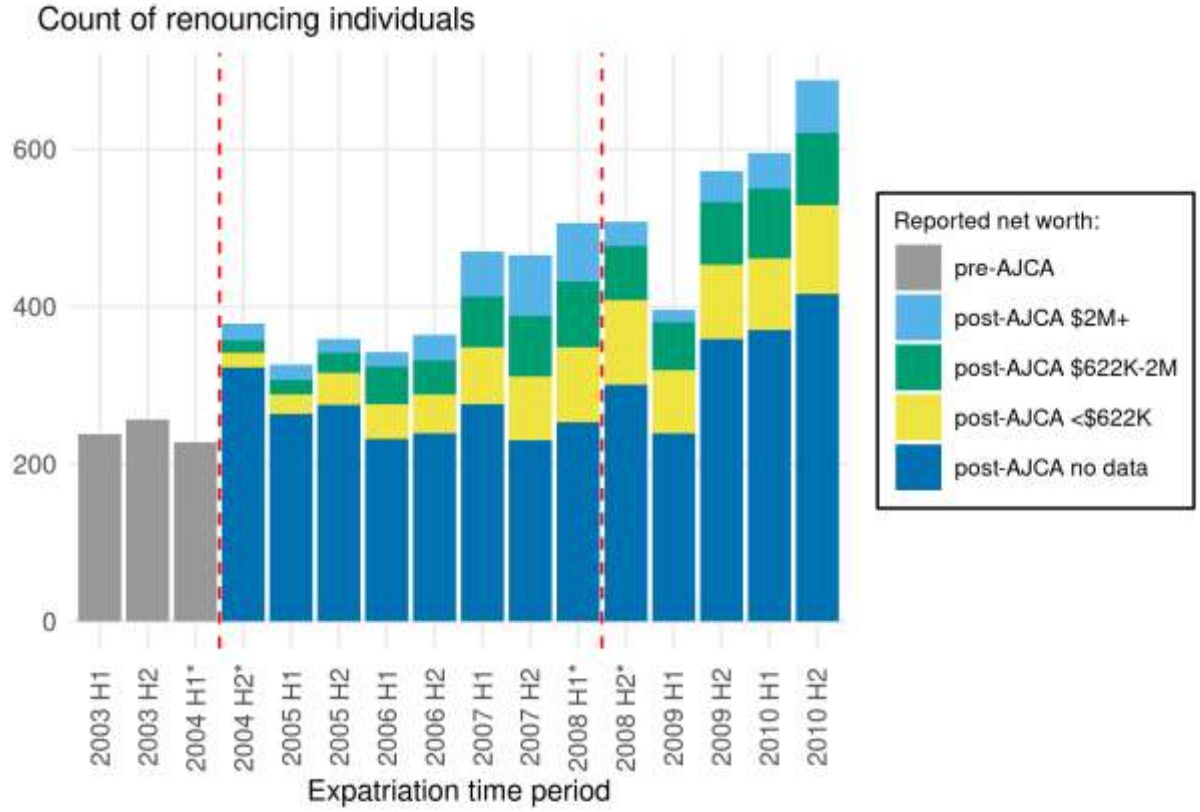

Notes: This figure plots the count of individuals renouncing in each half-year period grouped by reported net worth. For 2004, the periods are split around June 4, 2004, when the net worth threshold for designation as a covered expatriate increased from $622 thousand to $2 million, as part of the AJCA. For 2008, the periods are split around June 18, 2008, when the HEART Act's mark-to-market tax provisions went into place. For pre-AJCA expatriations, we can only observe whether an individual was above the net worth threshold (in gray) or not (in orange). Post-AJCA, most individuals filing Form 8854 report net worth, and thus can be grouped into three buckets based on the threshold changes, though there are still many individuals (in darker blue) without reported net worth data.

Focusing on the patterns around the 2008 HEART Act, some interesting patterns are visible. In the few periods prior to the HEART Act change, the number of renouncers reporting net worth of $622K-$2M and above $2M increased noticeably. This could reflect individuals accelerating their renunciations to avoid the mark-to-market tax, which was in discussion for at least several months prior to being passed and signed into law on, and affecting expatriations on or after, June 17, 2008.[35] The number of high-wealth renunciations fell in the second half of 2008 and first half

[35] Reichenberg Sherr (2008) notes that the expat provision ultimately passed as part of the HEART Act is "similar to…the expat provision in a prior bill, H.R. 3997, which was passed by both the House and Senate in December 2007 but did not get enacted due to other differences between the House and Senate bills."

of 2009, which again would be consistent with individuals moving renunciation forward to avoid the mark-to-market tax. However, the confounding effects of the financial crisis may also have affected the ability of U.S. citizens to move abroad or affected their decisions about whether to incur the costs of renunciation, and could also help to explain the drop in renunciations. Again, more complete information about these individuals would help to say more with greater certainty.

1.5.2.4. Responses to the net worth threshold

The data patterns discussed above suggest that some individuals responded to the changes in expatriation tax law. Further evidence of taxpayers responding to the expatriation tax rules can be seen by examining the pattern of filings with reported net worth above and below the $2M net worth threshold for designation as a covered expatriate.

As described above, expatriating individuals are subjected to a test that determines whether they are a covered expatriate. The test has three components, any one of which results in designation as a covered expatriate: (1) net worth above a threshold; (2) past-five-years average income tax liability above a threshold; and (3) failing to certify compliance on U.S. taxes for the five years prior to expatriation. Covered expatriate status results in additional filing requirements, as well as potential additional tax liability. Prior to the HEART Act in 2008, this tax liability was based on income earned during the 10 years following expatriation, which could be liable for U.S. income taxation. Since the HEART Act, this tax liability is a mark-to-market exit tax based on the value of all assets owned on the day prior to expatriation, with taxes applied to gains above a statutory exemption.

For most covered expatriates, the net worth threshold is the crucial component.[36] Since mid-2004, the net worth threshold has been constant at $2M. A histogram of renouncers' reported net worth around this threshold reveals a strong response, as shown in Figure 1.12: a sharp drop-off in the number of renouncers reporting net worth just above the threshold. This figure shows the aggregate histogram for all renunciations from mid-2004, when the AJCA took effect and net worth data become widely available, through 2018. Although not presented here for disclosure reasons, the pattern is also visible within each year.[37] There are several plausible explanations for this drop-off: some potential renouncers with net worth above $2M may have been discouraged from renouncing; some may have taken actions to reduce net worth below the $2M threshold (for example, by making gifts or charitable contributions); and some may have reported net worth lower than their actual net worth, in order to appear below the threshold. In addition, recall that only about half of renouncing individuals have a filed Form 8854 with reported net worth data available; it is possible that some individuals with net worth above the threshold chose not to file Form 8854.

[36] Among all covered expatriates, nearly 90% are over the net worth threshold, while only about 25% are over the average income tax liability threshold. The evidence suggests there is little direct response to the average income tax liability threshold, in that there is no bunching below the threshold (see Appendix Figure 1.24). One explanation is that it is harder for taxpayers to adjust an average based on past-5-years income tax liabilities than it is to adjust reported net worth at the point of expatriation.

[37] That the pattern is visible both before and after the HEART Act suggests that covered expatriate designation was viewed as costly even without the mark-to-market exit tax consequences introduced under the HEART Act.

Figure 1.12: Histogram of reported net worth around $2 million

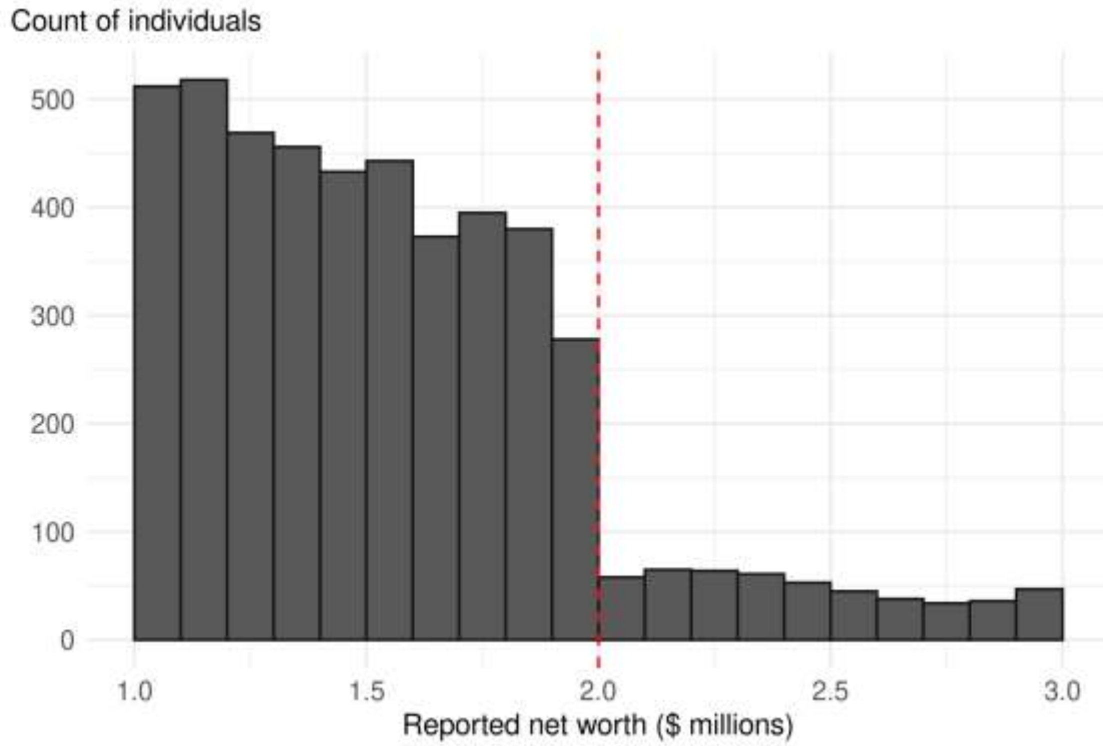

Notes: This figure plots the count of renouncers in each $100K bucket around the $2M threshold for designation as a covered expatriate. Renunciations with a filed Form 8854 and available reported net worth data, after the AJCA (mid-June 2004) through 2018, are included. The drop-off in filings with reported net worth occurs exactly at the cutoff for covered expatriate designation, suggesting it is this cutoff that is driving the observed behavior; there is no drop-off at either $1M or $3M, suggesting that round-number bunching can be ruled out as an explanation for the observed pattern (see Appendix Figure 1.23).

There is evidence that for some taxpayers, gifts may have been used to get below the threshold. 8% of the individuals who report net worth of $1-2M would have had net worth above $2M if gifts they reported making in the 0-2 years prior to renunciation were added to their reported net worth. A handful of individuals similarly would move from below the threshold to above it if their pre-renunciation charitable contributions were added to their reported net worth.[38] Still, even after adjusting the reported net worth amounts to include recent gifts and charitable contributions, a large "hole" to the right of the threshold remains. One feature of Form 8854 (the expatriation tax

[38] For this analysis, I rely on gift amounts as reported on Form 709 and charitable contributions reported on Schedule A.

form) is that it requires individuals to provide a balance sheet with assets listed by asset type; although not presently available, these data could in future be used to further explore the patterns shown here.

In sum, there is evidence consistent with the hypothesis that renunciations were responsive to these tax law changes, although it is not possible to draw conclusions with certainty. When pairing the patterns shown here with the income and wealth trends discussed earlier in Section 1.4.3, the strongest trends seem to be that a group of high-wealth and high-income individuals chose to renounce after the AJCA and before the HEART Act. If these renunciations were tax-motivated and made in anticipation of the mark-to-market exit tax, this suggests that the exit tax was perceived by many taxpayers as costly and worth avoiding (a view further supported by the strong and observable behavioral response to the net worth threshold).

Although it is not possible to give a single answer to the question "Why are they renouncing?", the preceding analyses help to provide some resolution. The results suggest that the recent increase in renunciations was caused by increased compliance costs for those already living abroad, and that some individuals' renunciation decisions during the mid-2000s were at least in part a response to changes, or expectations of changes, to the expatriation tax system.

1.6. Policy consequences

Building on the findings above about who is renouncing and why, this section answers my third and final research question: What are the policy consequences? I first consider the revenue impacts of recent renunciations. I then discuss what lessons can be learned from the policy changes over the last two decades and conclude by discussing what these findings suggest about the value of U.S. citizenship.

1.6.1. Revenue impacts

Considering the revenue impacts of recent renunciations, the evidence suggests that, although most do not have any effect, or at most a small one, a handful of renunciations by very high-wealth and high-income individuals could have substantial revenue impacts.

A simple way to think about the direct revenue impacts of renunciations is to consider the tax liabilities renouncers had in the years leading up to their renunciation and assume that these liabilities would have continued had they not renounced. Focusing on the year just prior to renunciation, Figure 1.13 shows that the share of renouncers with no tax liability in the year prior to renunciation has increased markedly; since 2013, about two-thirds of renouncers linked to a Form 1040 filing as the primary filer in the year prior to renunciation had no liability on that return. These linked returns represent between half and two-thirds of all renouncers; the remainder are mostly those without TINs, or with TINs but no linked filings, who also likely had no U.S. tax liability. Including them in the proportions would further increase the share of renunciations with no revenue impact.[39]

[39] Appendix Figure 1.22 shows the annual counts, including those without filings as one group, and splitting those with filings into the three liability buckets used here. Most of the recent increase in renunciations is from individuals with no filings or no liability.

Figure 1.13: Pre-renunciation tax liability

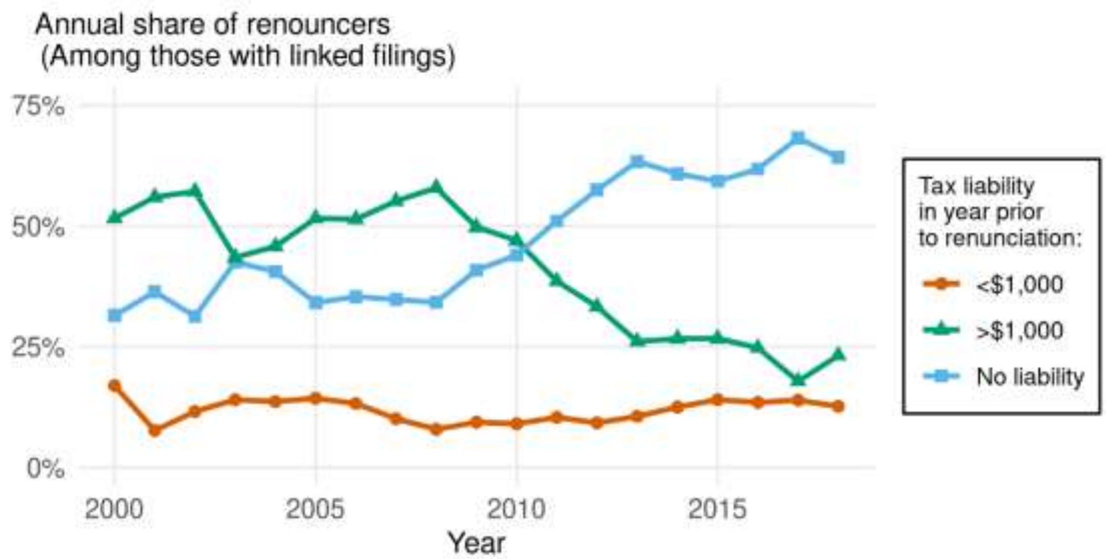

Notes: This figure reports the share of individuals renouncing each year with a pre-renunciation tax liability of zero, <$1,000, or >$1,000, among those who are linked as a primary filer on a Form 1040.

Although many, indeed most, renunciations probably have a negligible revenue impact, this is not universally true, nor is it necessarily the case that this pattern will hold indefinitely. As shown earlier in Section 1.4.3, a handful of wealthy and high-income renouncers can have an outsize impact on the average net worth and income of those renouncing, and thus on the estimated revenue impacts. If policymakers are concerned about renunciation purely from a revenue perspective, the wealthy and high-income are where their focus should continue to be. The experience of the past two decades does provide some evidence that policy can help discourage renunciation by these individuals. The prevalence of especially wealthy and high-income individuals among those renouncing between 2004 and 2008 suggests that the introduction of the mark-to-market tax was perceived as costly, and thus may have had some success in discouraging subsequent high wealth and income taxpayers from renouncing (although unable to stop those who could renounce before its enactment, an issue I discuss below). In addition, for those still choosing to renounce, the mark-to-market tax helps to mitigate the revenue impact. One high-profile example of this is the renunciation of U.S. citizenship by Facebook co-founder Eduardo Saverin.[40] Although his renunciation meant the U.S. lost out on future income and estate tax revenue, this was at least partly offset by his exit tax liability (which, according to reports in the public press, likely was in the hundreds of millions of dollars (Benoit 2012).

Table 1.17 in Appendix A provides an additional set of summary statistics that support these conclusions about the revenue impacts of recent renunciations. More than half of Droppers had no liability during all five years prior to their renunciation. Movers are more likely to have had non-zero liabilities, but for most individuals, these are still relatively small. The median non-zero liability for Movers in the year prior to renunciation was about $12K, or $8K when considering the average over the five years prior to renunciation. These median values are about 10 times smaller than the mean values, again illustrating that a few outliers have a large impact while most individuals do not.

Of course, liabilities can change from year to year and assuming that they would stay constant may not always be correct. A more refined estimate of the revenue impacts of renunciation could take several routes. To get a more precise estimate of the direct revenue impacts, one could more carefully forecast what the path of tax liabilities would have been, absent renunciation. This could consider the path of liabilities prior to renunciation, as well as the renouncer's age and assumptions about retirement age and life expectancy. In addition, the revenue impacts should include estimated effects on future estate tax liabilities, and the revenue raised from expatriation tax liabilities of covered expatriates. Still, even taking account of these refinements the conclusion is unlikely to change: most renunciations have had minimal revenue impact, but a handful probably had a significant impact.

[40] According to the quarterly publication of expatriating individuals in the Federal Register, Saverin renounced his U.S. citizenship in the first quarter of 2012, prior to Facebook, Inc's IPO on May 18, 2012 (77 FR 25538).

A secondary impact on revenue could come from "brain drain". Academic literature on this topic has focused mainly on high-skilled migration from less-developed to more-developed countries (Gibson and McKenzie 2011), but in principle it could also matter for the U.S. It is possible that some Movers' renunciations could generate negative spillover effects in the U.S., for example if Movers close or relocate U.S. businesses when they move abroad, or postpone entrepreneurial activity and innovation until after renunciation. Given the small number of Movers to date, this is unlikely to have had a significant impact, but in principle could become relevant if future policies led to increasing numbers of Movers. This is less of a concern when considering the impacts of renunciation by Droppers, as their economic activity is probably concentrated abroad and thus their renunciations are unlikely to have spillover effects in the U.S.

Finally, consider that the incentives affecting the outflow of citizens (renunciation) should also affect the inflow of citizens (immigration and naturalization). The academic literature on immigration points to economic incentives as one factor determining whether, when, and where individuals migrate (Freeman 2006). As shown above, expatriation tax rules did affect renunciation decisions by some U.S. citizens on the margin, particularly the wealthy and high-income. It is plausible that these rules, and the tax costs and benefits of U.S. citizenship, would similarly affect the decisions of those considering in-migration to the U.S. Mason (2016) raises the concern that citizenship taxation could discourage marginal wealthy or high-income migrants; Kim (2017) disagrees, arguing that it is U.S. immigration law, not tax law, that is the real obstacle for highly skilled and educated immigrants. The key determinant of the importance of the tax law effect is the existence of at least some individuals considering in-migration who are on the margin. If the distribution of those considering in-migration is comparable to those considering renunciation, then U.S. tax law could discourage some individuals on the margin from migrating to the U.S. or naturalizing once in the U.S. Given the relative magnitudes (for the U.S., naturalizations are two orders of magnitude higher than renunciations, as I discuss below), this could have significant implications for U.S. tax revenue and economic activity.

1.6.2. Policy lessons

Studying the renunciation responses to recent tax policy changes reveals two additional lessons. First, unintended side effects matter: FATCA appears to have induced some U.S. citizens abroad to renounce citizenship, and the resulting social cost should be considered when evaluating FATCA. Second, timing matters: the timing of the AJCA and HEART Act legislation may have allowed some high-wealth individuals to renounce in advance of the exit tax taking effect.

The analysis in Section 1.5.1 showed that the increase in renunciations in the last decade was in part an unintended side effect of FATCA and other related policies that, while having some positive revenue impacts, imposed additional compliance costs on those maintaining financial accounts abroad. Does the U.S. value those foreign-resident U.S. citizens? I argue that the answer is yes. It may at first seem that these individuals' welfare should be discounted; because such individuals are often called "accidental Americans" one might think their renunciations do not have a social cost. The treatment of citizenship under U.S. nationality law, however, suggests this is not the case. In principle, the U.S. could further restrict citizenship but so far has not. This reveals that the U.S. indeed places some value or social benefit on maintaining citizenship for these individuals.[41] Thus, the U.S. *loses* value, or experiences a social cost, when those abroad renounce citizenship, and this cost should be included when evaluating the overall effects of FATCA.

---

[41] My thanks to Dhammika Dharmapala for a helpful discussion about this topic.

The experience of renunciations during the 2000s also illustrates the importance of policy timing, and how anticipatory action can partially negate some of the intended effects of legislation. As shown above, the years between the AJCA and HEART Act saw a handful of wealthy individuals renouncing citizenship, perhaps influenced in part by a desire to renounce before the imposition of the mark-to-market exit tax which was being discussed but not yet implemented. A resulting lesson is thus that the speed of debate and implementation becomes more important when considering a policy that is intended to target a small group of people who are sophisticated and well-informed about potential policy changes.

1.6.3. The value of U.S. citizenship

Finally, it is important to put recent renunciations in context. This paper focuses on those dropping citizenship, motivated by the recent increase in renunciations. However, to evaluate the effect of the tax system on citizenship decisions overall, consider the rest of the picture: most individuals do *not* choose to renounce citizenship, and there is also a large number each year *gaining* U.S. citizenship.

Consider first the naturalizations: although the relative increase in renunciations over the last decade is remarkable, the net flow (naturalizations less renunciations) is still vastly tilted towards in-migration. Figure 1.14 plots the annual count of naturalizations (those receiving U.S. citizenship) in gray, and renunciations (plotted with negative values) in orange. The renunciations are just barely distinguishable at the bottom of the graph, two orders of magnitude smaller than the naturalizations. In every year between 1998 and 2018, the number of naturalizations was above

400,000. This compares to a *total* of roughly 40,000 citizenship renunciations between 1998 and 2018.[42]

Figure 1.14: Naturalizations vs. citizenship renunciations

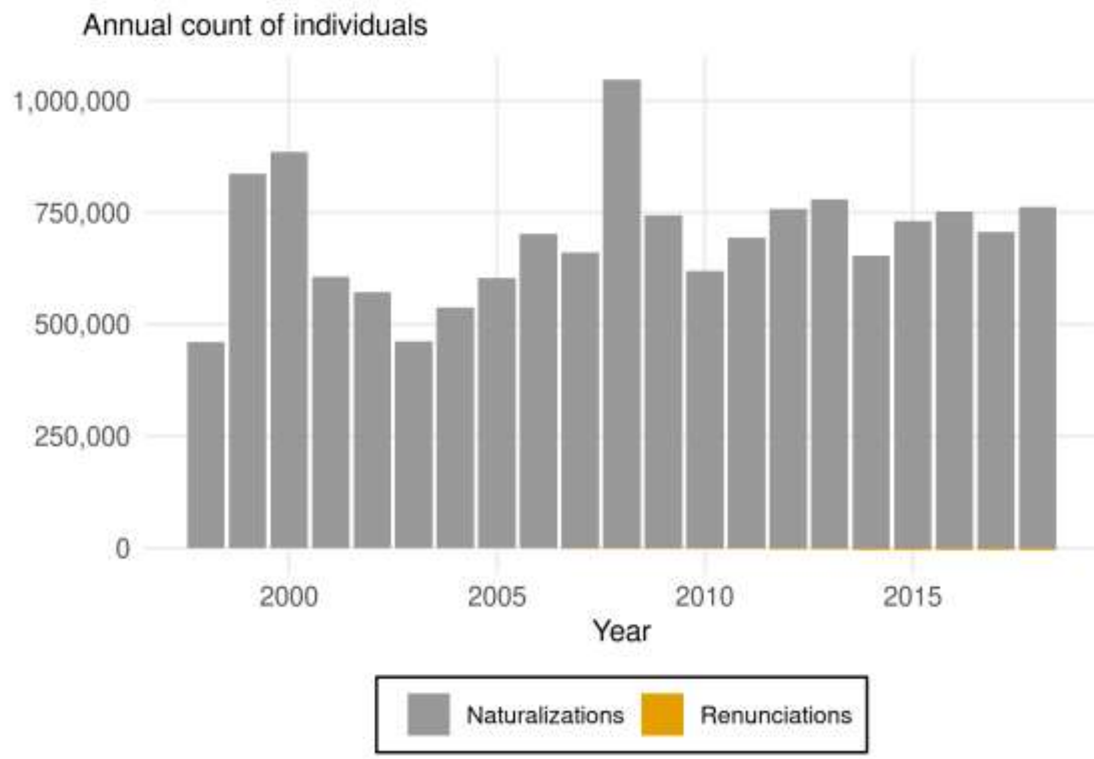

Notes: This figure plots in gray the total number of U.S. naturalizations each year from 1998-2018, from DHS, 2019 Yearbook of Immigration Statistics, Table 20; in orange with negative values are the annual counts of those renouncing citizenship.

What about those who already have citizenship, and choose not to renounce it? This describes almost all U.S. citizens. There are more than 300 million such individuals, and typically fewer than 5,000 renouncing each year. The number of renunciations is still tiny even when compared to the stock of U.S. citizens abroad, who could more readily renounce. Although the exact number of U.S. citizens living abroad is not known, some estimates put it at perhaps nine million, and the number filing taxes from foreign addresses is more than one million per year. A few thousand

[42] As noted above, due to data accessibility I focus in this paper on citizenship and not long-term residency status, but similar arguments can be made for the long-term resident population, with similar conclusions about the effect of the tax system on individuals' decisions. In each year, the number of individuals relinquishing long-term residency status is far lower than the number applying for it.

renunciations per year thus represents, as a conservative upper bound, less than half of 1 percent of those living abroad.[43] This suggests another lesson from the fact that the increased compliance costs under FATCA induced some individuals abroad to drop their U.S. citizenship: those costs did *not* induce vastly many more foreign-resident U.S. citizens to drop citizenship, implying that for those individuals the maintenance of U.S. citizenship was worth incurring the resulting financial and hassle costs of complying with new regulations, and thus that they place a relatively high value on U.S. citizenship.

1.7. Conclusions

Because the U.S. tax system applies to its citizens' worldwide income and estates, citizenship and taxes are more closely connected for the U.S. than for nearly any other country. Using novel administrative tax microdata on the population of individuals who have dropped U.S. citizenship over the past twenty years, this paper demonstrates that this connection can have substantial impacts on taxpayer behavior, including the decision to maintain or renounce citizenship.

The preceding analyses provide a detailed understanding of who is renouncing and why. The recent increase in renunciations has come mainly from those who have long filed U.S. taxes from abroad – that is, mainly from Droppers, not Movers. These Droppers' renunciations were primarily an unintended side effect of the increased compliance costs brought on by FATCA and other offshore financial enforcement during the 2010s. And although renouncers on average are wealthier and higher-income than the U.S. population, most recent renouncers had low or zero pre-

[43] I am not the first to draw this comparison; a similar point was made by Elise Bean in her testimony before the House Subcommittee on Government Operations in a hearing titled "Reviewing the Unintended Consequences of the Foreign Account Tax Compliance Act," held on April 26, 2017. In some respects, the discussion of renunciations is similar to that of corporate inversions: although the absolute number occurring is relatively small, there is still significant public press and legislative focus on the issue.

renunciation U.S. tax liability, suggesting that their renunciations may not have a significant revenue impact.

The evidence reveals that citizenship decisions are connected to U.S. tax policy, most notably that the compliance costs of increased offshore enforcement may have led thousands of U.S. citizens abroad to drop their citizenship. For these individuals the costs of renunciation, both financial and emotional, surely were quite large. Still, the total number of renunciations remains relatively small, whether compared to estimates of the remaining population of U.S. citizens living abroad, filing taxes from foreign addresses, or newly gaining U.S. citizenship, and in purely financial terms, the revenue impact of their renunciations is likely to be small. All this together suggests that U.S. citizenship has historically been perceived as valuable by most who hold it, and remains so today.

That citizenship decisions and the tax system are connected should be accounted for when considering changes to the tax system. The attractiveness of citizenship renunciation depends crucially on the current tax system as well as expectations about its future, relative to alternative foreign tax systems. Individuals determine the expected costs and benefits of retaining or dropping citizenship, factoring in the potential for future tax increases (or decreases) or even entirely new taxes, such as an annual wealth tax. This determination may be particularly relevant for younger individuals facing a future stream of annual tax liabilities, for entrepreneurs considering the potential future net-of-tax gains to their innovation, and for the wealthy considering potential future estate tax liabilities. Those considering moving to the U.S., or naturalizing as U.S. citizens, may also be influenced by the tax system. Policymakers should not ignore citizenship renunciation and naturalization as potentially important margins of response.

1.8. Appendix A: Additional figures and tables

Figure 1.15: Comparison of income for renouncers, foreign filers, and all tax filers, mean and median

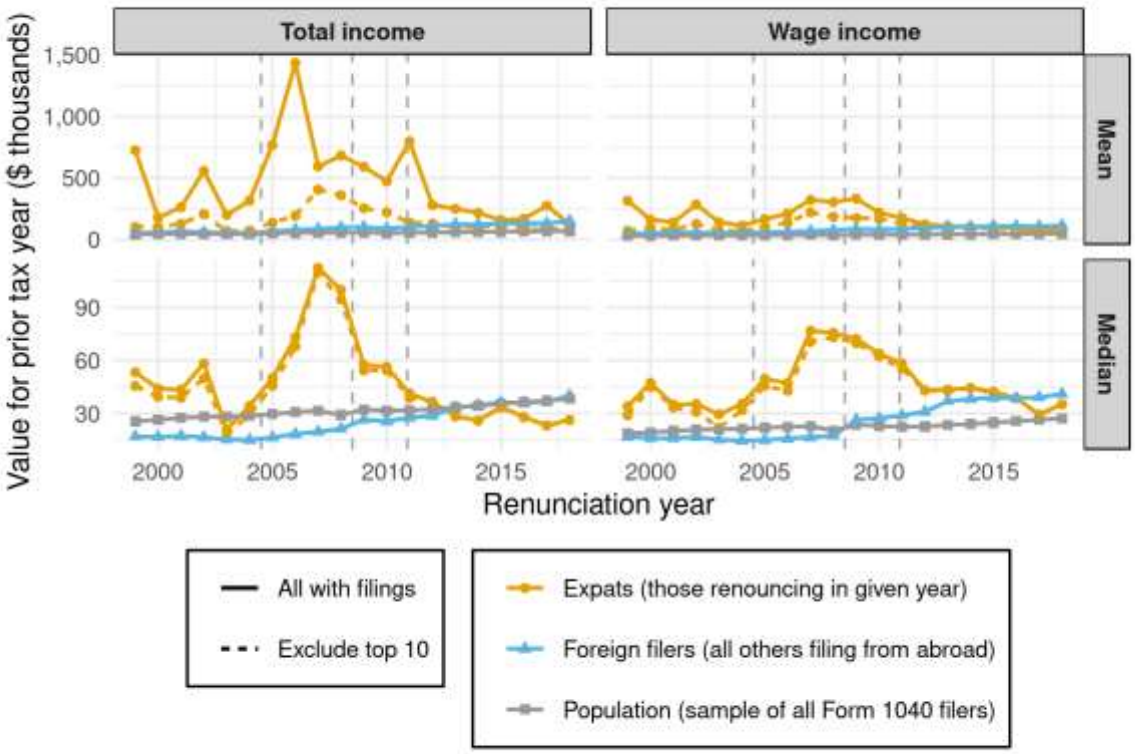

Notes: This figure compares the income of renouncers in the year prior to renunciation to two comparison groups: all other foreign filings, and a sample of the population of Form 1040 filing. For renouncers, only primary filers with linked filings are included. The three vertical dashed lines represent three key dates related to expatriation tax law: the 2004 AJCA, the 2008 HEART Act, and 2010 FATCA. The solid line includes all individuals; the dashed line removes the top 10 in each year.

Figure 1.16: Comparison of income among renouncers, Movers vs. Droppers, mean and median

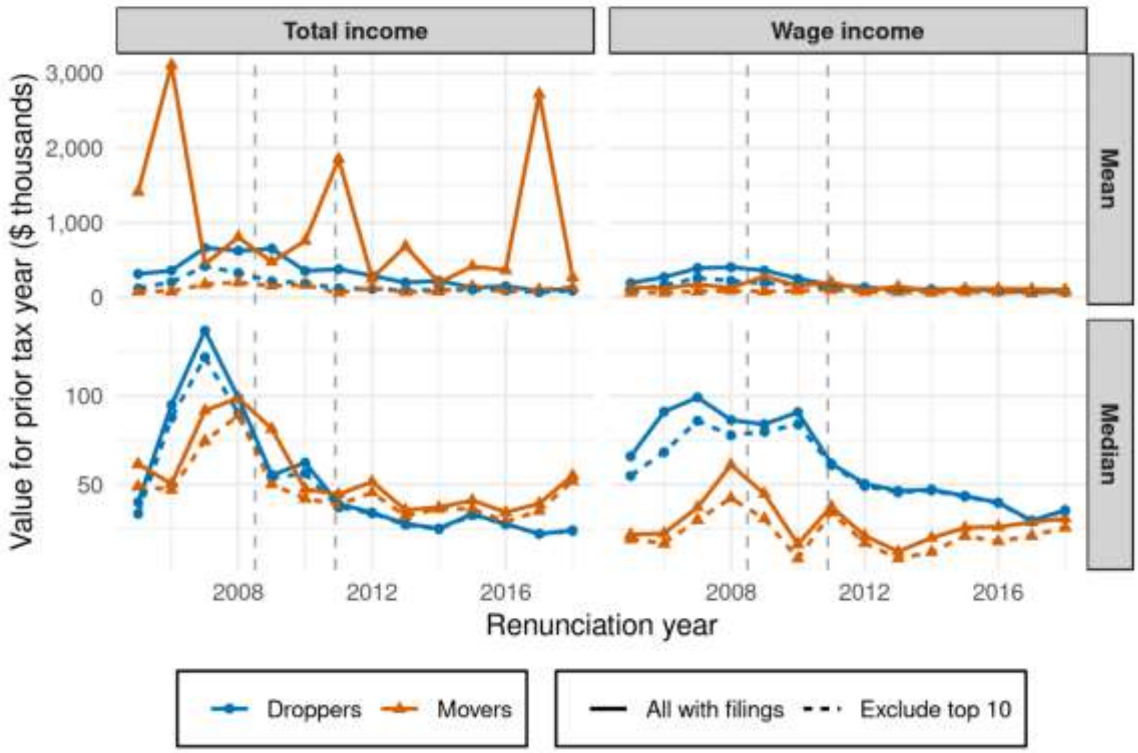

Notes: This figure comparison the income in the year prior to renunciation for renouncers with linked Form 1040 filings as primary filers. The mean values are calculated separately among Movers and Droppers. Renouncers with no filings or no TINs are excluded.

Figure 1.17: Comparison of reported net worth among renouncers, Movers vs. Droppers, mean and median

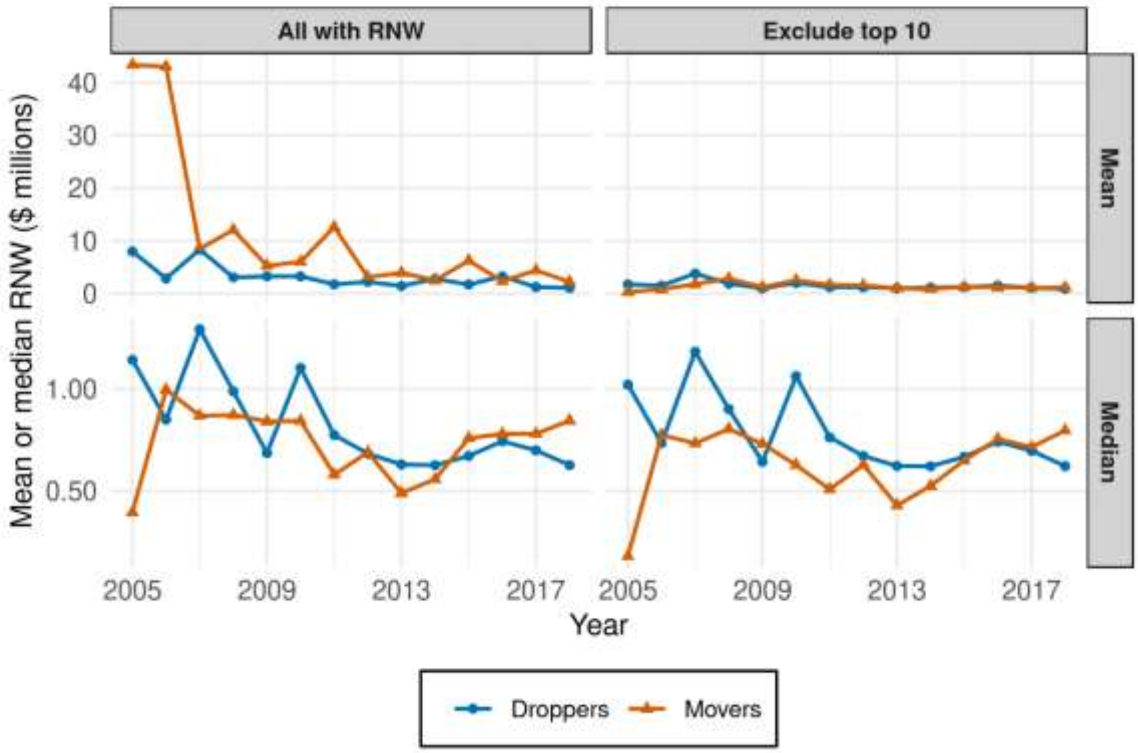

Notes: This figure compares reported net worth among those renouncing each year, separately for Movers and Droppers. Only those with reported net worth data available are included. The left panel includes all Movers and Droppers; right panel drops the top 10 Movers and Droppers, by reported net worth, each year.

Table 1.7: Top destinations, split by renouncer reported net worth

		Renouncers, split by reported net worth				
Rank	**All renouncers**	**$2M+**	**$622K-2M**	**$0-622K**	**No RNW**	**No 8854**
1	Canada	Canada	Canada	Canada	Canada	Canada
2	Switzerland	United Kingdom	Switzerland	Switzerland	Switzerland	Switzerland
3	United Kingdom	Switzerland	United Kingdom	United Kingdom	United Kingdom	United Kingdom
4	Germany	Hong Kong	Hong Kong	Hong Kong	Germany	Germany
5	Hong Kong	Australia	Australia	Germany	South Korea	South Korea
6	Australia	France	Germany	Australia	China	Singapore
7	South Korea	Germany	France	Netherlands	Norway	Hong Kong
8	Singapore	Singapore	Singapore	Taiwan	Hong Kong	Australia
9	Taiwan	Taiwan	Taiwan	France	France	Taiwan
10	France	China	Belgium	Singapore	*Japan*	Belgium

Notes: This table presents the top 10 destinations among all renouncers, and then within each reported net worth group. All renunciations between 2005 and 2018 are included.

Table 1.8: Top destinations, split by renouncer classification

		Renouncers, split by classification			
Rank	**All renouncers**	**Mover**	**Dropper**	**No Filings**	**No TIN**
1	Canada	Canada	Canada	Switzerland	Switzerland
2	Switzerland	United Kingdom	Switzerland	Canada	Canada
3	United Kingdom	Switzerland	United Kingdom	United Kingdom	Germany
4	Germany	Hong Kong	Hong Kong	Germany	United Kingdom
5	Hong Kong	Taiwan	Australia	South Korea	South Korea
6	Australia	Germany	Germany	Hong Kong	Singapore
7	South Korea	South Korea	France	Australia	Hong Kong
8	Singapore	China	Netherlands	France	Taiwan
9	Taiwan	Australia	Singapore	Singapore	Belgium
10	France	Singapore	New Zealand	Taiwan	Australia

Notes: This table presents the top 10 destinations among all renouncers, and then within each reported classification. All renunciations between 2005 and 2018 are included.

Table 1.9: Top destinations, by count and by share of unique U.S. tax filers

Total renunciations 1998-2018			Total renunciations 2008-2018			Renunciations / Unique US tax filers Renunciations '08-'18; Filers '07-'17		
Rank	**Jurisdiction**	**Count**	**Rank**	**Jurisdiction**	**Count**	**Rank**	**Jurisdiction**	**Share**
1	Canada	8,970	1	Canada	8,350	1	Liechtenstein	34.1%
2	Switzerland	6,330	2	Switzerland	5,900	2	Switzerland	17.2%
3	United Kingdom	4,270	3	United Kingdom	3,480	3	Monaco	13.9%
4	Germany	2,790	4	Germany	1,950	4	Belgium	7.0%
5	Hong Kong	1,880	5	Hong Kong	1,630	5	Hong Kong	5.5%
6	Australia	1,410	6	Australia	1,240	6	Kuwait	5.4%
7	South Korea	1,310	7	Singapore	1,070	7	Singapore	5.4%
8	Singapore	1,180	8	South Korea	890	8	Panama	5.0%
9	Taiwan	930	9	Taiwan	850	9	St. Kitts and Nevis	4.8%
10	France	820	10	France	730	10	Somalia	4.8%

Notes: this table presents the top 10 destinations by total count (left two panels, for 1998-2018 and 2008-2018). The third panel shows the number of renunciations as a share of the unique U.S. tax filers from that jurisdiction, counting all renunciations from 2008-2018 and dividing by the number of unique U.S. tax filers from 2007-2017.

Table 1.10: Summary statistics for individual-level regression data

Variable	N	Mean	St. Dev.	Min	Pctl(25)	Median	Pctl(75)	Max
Binary outcome (100/0)								
Renounce in following year	4,831,000	0.343	5.849	0	0	0	0	100
Covariates								
Total Positive Income ($M)	4,831,000	0.14	1.48	0	0.02	0.05	0.11	1,200
Wage share of TPI (%)	4,831,000	0.64	0.44	0.00	0.00	0.95	1.00	1.00
Age (years)	4,790,000	48	17	16	34	46	60	100
Binary covariates (1/0)								
Positive tax liability	4,831,000	0.38	0.49	0	0	0	1	1
Had any Sch C income	4,831,000	0.12	0.32	0	0	0	0	1
Had any Sch E income	4,831,000	0.14	0.34	0	0	0	0	1
Received an IRS notice	4,831,000	0.15	0.35	0	0	0	0	1
Made Sch A charity deduction	4,831,000	0.07	0.25	0	0	0	0	1
Filed a gift tax form	4,831,000	0.00	0.05	0	0	0	0	1
Diff-in-diff covariates (1/0)								
IGA jurisdiction								
as of 2019	4,831,000	0.932	0.25	0	1	1	1	1
as of 2017	4,831,000	0.930	0.26	0	1	1	1	1
as of 2015	4,831,000	0.873	0.33	0	1	1	1	1
Tax haven jurisdiction	4,831,000	0.082	0.27	0	0	0	0	1
Post (tax years >= 2010)	4,831,000	0.766	0.42	0	1	1	1	1

Notes: This table reports the summary statistics for the individual-year level data used in regression analysis with results shown in Table 1.4. The population is all Form 1040 filings from foreign addresses for tax years 2007-2017, excluding any Mover renouncers. Values are rounded for disclosure purposes.

Table 1.11: *Individual regression results, including all linked renouncers*

Dependent variable:	Binary: Renounce in following year (100/0)					
	[1]	[2]	[3]	[4]	[5]	[6]
Total Positive Income	0.0171	0.017	0.0176	0.0137	0.0140***	0.0137
($ millions)	(0.0105)	(0.0104)	(0.0105)	(0.0098)	(0.0018)	(0.0097)
Wage share	-0.0934	-0.071	-0.1216	-0.1345**	-0.1293***	-0.1250***
(% of TPI)	(0.0804)	(0.0614)	(0.0748)	(0.0558)	(0.0077)	(0.0385)
Positive tax liability	-0.0692	-0.0717	-0.0579	-0.0632	-0.0562***	-0.0546
(1/0)	(0.0655)	(0.0665)	(0.0644)	(0.0472)	(0.0061)	(0.0453)
Any Sch C income	0.0843***	0.0874***	0.0624**	0.0570**	0.0433***	0.0469**
(1/0)	(0.0278)	(0.0309)	(0.0254)	(0.0196)	(0.0091)	(0.0195)
Any Sch E income	0.0127	0.0105	-0.0046	0.018	-0.0011	0.0077
(1/0)	(0.0381)	(0.0392)	(0.0376)	(0.0270)	(0.0083)	(0.0264)
Schedule A charity	-0.0052	-0.0131	0.0264	-0.1158	-0.0928***	-0.0907
(1/0)	(0.0915)	(0.0894)	(0.0938)	(0.0972)	(0.0115)	(0.0975)
Filed gift tax return	2.5150***	2.5109***	2.4779***	2.4555***	2.4139***	2.4046***
(1/0)	(0.4208)	(0.4180)	(0.4150)	(0.4284)	(0.0568)	(0.4188)
Received any notice	0.1163**	0.1123**	0.0376*	0.1241**	0.0454***	0.0506*
(1/0)	(0.0560)	(0.0563)	(0.0197)	(0.0543)	(0.0082)	(0.0279)
Age		0.0016			0.0020***	0.0019
(years)		(0.0013)			(0.0002)	(0.0012)
Constant	0.4168***	0.3304***				
	(0.1370)	(0.0720)				
Year FE	No	No	Yes	No	Yes	No
Jurisdiction FE	No	No	No	Yes	Yes	No
YearXJurisdiction FE	No	No	No	No	No	Yes
Observations	4,835,000	4,793,000	4,835,000	4,835,000	4,793,000	4,793,000
Adjusted R^2	0.0006	0.0006	0.0014	0.0046	0.0054	0.0067
Mean dep. var.	0.343	0.343	0.343	0.343	0.343	0.343

Notes: *p<0.1; **p<0.05; ***p<0.01. Standard errors, clustered by year and by jurisdiction, are shown in parentheses. The dependent variable is coded as 100 or 0 so that the coefficient estimates represent the effect in percentage points for each covariate, holding all others constant. All expatriates that can be linked to Form 1040 filings as a primary filer are included.

Figure 1.18: Scaled coefficient estimates for selected covariates, individual LPM

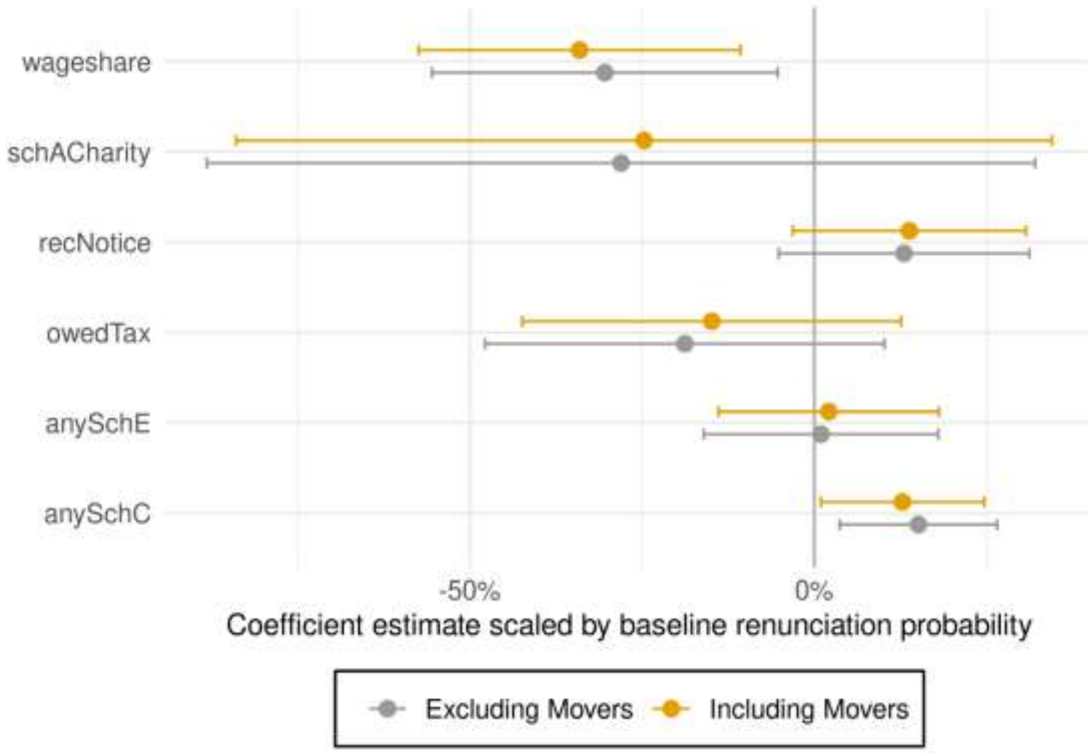

Notes: This figure plots the coefficient estimates for selected covariates of the fully saturated linear probability model reported in column 6 of Table 1.4 (excluding Movers) and Table 1.11 (including Movers). The estimates are scaled by the mean dependent variable to show the estimated effect on the probability of a given type of expatriation, in percent. 95% confidence intervals using standard errors clustered by year and jurisdiction are shown around each point estimate. The covariate for filing a gift tax form is excluded from this figure because it dominates the others, with a scaled point estimate suggesting it is associated with a more than 500% increase in the probability of renunciation.

Table 1.12: Summary statistics for jurisdiction-level regression data

Variable	N	Mean	St. Dev.	Min	Pctl(25)	Median	Pctl(75)	Max
Outcome								
Renunciation share	213	0.013	0.030	0	0	0.005	0.0128	0.3415
Covariates								
Tax haven	213	0.174	0.380	0	0	0	0	1
Citizenship-for-sale	213	0.047	0.212	0	0	0	0	1
English-speaking	213	0.094	0.292	0	0	0	0	1
WGI Rule of Law	205	0.490	0.284	0	0.251	0.474	0.728	0.996
WGI Political Stability	205	0.497	0.286	0.012	0.255	0.493	0.769	0.985
Passport value	197	0.469	0.297	0.015	0.214	0.415	0.733	0.979
Average change in RGDP	188	3.218	2.370	-4.609	1.600	3.177	4.766	11.509

Notes: This table reports the summary statistics for the jurisdiction-level data used in the regression analysis with results shown in Table 1.5. The population is all jurisdictions with any U.S. tax filings for tax years 2007-2017. Certain covariates are only available for a subset of jurisdictions. I define citizenship-for-sale jurisdictions as those that began such a program prior to 2017: Antigua and Barbuda, Bulgaria, Comoros, Dominica, Grenada, Malta, St. Kitts and Nevis, Serbia, St. Lucia, and Vanuatu (Christians 2017).

Table 1.13: Individual-level regression results with jurisdiction-level covariates

Dependent variable:	Binary: Renounce in following year (100/0)					
	[1]	[2]	[3]	[4]	[5]	[6]
Total Positive Income	0.0102	0.0068	0.0073	0.0108	0.0074	0.0079
($ millions)	(0.0063)	(0.0060)	(0.0063)	(0.0063)	(0.0073)	(0.0072)
Wage share	-0.055	-0.0922**	-0.0961**	-0.073	-0.1093**	-0.1122**
(% of TPI)	(0.0603)	(0.0422)	(0.0418)	(0.0567)	(0.0389)	(0.0389)
Positive tax liability	-0.0834	-0.0869*	-0.0883*	-0.0727	-0.0772	-0.0784
(1/0)	(0.0650)	(0.0465)	(0.0462)	(0.0638)	(0.0442)	(0.0445)
Any Sch C income	0.0937***	0.0864***	0.0887***	0.0749**	0.0691***	0.0710***
(1/0)	(0.0278)	(0.0236)	(0.0221)	(0.0274)	(0.0199)	(0.0198)
Any Sch E income	0.0042	-0.0044	-0.0044	-0.0129	-0.0208	-0.0207
(1/0)	(0.0383)	(0.0354)	(0.0348)	(0.0377)	(0.0342)	(0.0338)
Schedule A charity	-0.0218	-0.114	-0.1148	0.0084	-0.0828	-0.0834
(1/0)	(0.0856)	(0.0912)	(0.0928)	(0.0874)	(0.0930)	(0.0937)
Filed gift tax return	2.2807***	2.2270***	2.2497***	2.2419***	2.1897***	2.2117***
(1/0)	(0.4172)	(0.4222)	(0.4266)	(0.4108)	(0.4141)	(0.4183)
Received any notice	0.1023*	0.1066*	0.1055*	0.025	0.032	0.0311
(1/0)	(0.0547)	(0.0569)	(0.0569)	(0.0191)	(0.0283)	(0.0286)
Age	0.0019	0.0020*	0.0019*	0.0023*	0.0024**	0.0023**
(years)	(0.0013)	(0.0011)	(0.0011)	(0.0012)	(0.0010)	(0.0010)
Tax haven		0.7891**	0.8146**		0.7949*	0.8255**
(1/0)		(0.3659)	(0.3698)		(0.3657)	(0.3698)
Citizenship-for-sale		-0.2167	-0.2056		-0.2018	-0.1946
(1/0)		(0.1338)	(0.1445)		(0.1373)	(0.1484)
English-speaking		0.2075***	0.2124***		0.2043***	0.2109***
(1/0		(0.0435)	(0.0453)		(0.0568)	(0.0572)
Rule of Law index		0.2542***	0.4921***		0.2257***	0.4409**
(percentile)		(0.0672)	(0.1799)		(0.0655)	(0.1752)
Political Stability index		0.2708***	0.2558***		0.2842***	0.2603**
(percentile)		(0.0775)	(0.0755)		(0.0854)	(0.0877)
Passport ranking			-0.3995*			-0.3806
(percentile)			(0.2280)			(0.2226)
Change in Real GDP			-0.035			-0.0374
(percentage points)			(0.0255)			(0.0253)
Constant	0.2890***	-0.2045**	0.0171			
	(0.0712)	(0.0810)	(0.1428)			
Year FE	No	No	No	Yes	Yes	Yes
Observations	4,790,000	4,786,000	4,752,000	4,790,000	4,786,000	4,752,000
Adjusted R^2	0.0005	0.0030	0.0030	0.0014	0.0038	0.0039
Mean dep. var.	0.343	0.343	0.343	0.343	0.343	0.343

Notes: *p<0.1; **p<0.05; ***p<0.01. Standard errors clustered by year are shown in parentheses. Jurisdiction-level covariates are averages over the full time period, and not time-varying. Observations are individual-year.

Table 1.14: Jurisdiction-level regression results, CFS robustness checks

	Dep. var.: Total renunciations/unique foreign filers							
	Robustness 1: Remove CFS covariate				*Robustness 2: Include Haven X CFS interaction*			
	[1]	[2]	[3]	[4]	[5]	[6]	[7]	[8]
Tax haven	0.0253**	0.0254**	0.0332**	0.0160*	0.0267**	0.0273**	0.0367**	0.0172*
(1/0)	(0.0121)	(0.0123)	(0.0147)	(0.0083)	(0.0131)	(0.0137)	(0.0168)	(0.0096)
Citizenship-for-sale					-0.0054*	-0.0070***	-0.0085***	-0.0065***
(1/0)					(0.0028)	(0.0024)	(0.0031)	(0.0025)
Tax haven X CFS					-0.0077	-0.0081	-0.0139	-0.0006
(1/0)					(0.0118)	(0.0117)	(0.0145)	(0.0103)
English-speaking	-0.0074	-0.0144	-0.0189*	-0.0064	-0.0049	-0.0111	-0.014	-0.0051
(1/0	(0.0080)	(0.0097)	(0.0111)	(0.0065)	(0.0067)	(0.0079)	(0.0088)	(0.0055)
Rule of Law index		0.0328***	0.0194***	0.0206***		0.0323***	0.0186**	0.0201***
(percentile)		(0.0093)	(0.0072)	(0.0063)		(0.0093)	(0.0074)	(0.0063)
Political Stability index		-0.0036	0.0002	-0.0049		-0.0034	0.0014	-0.0042
(percentile)		(0.0067)	(0.0074)	(0.0050)		(0.0070)	(0.0078)	(0.0051)
Passport ranking			0.0120**	0.0089			0.0112**	0.0087
(percentile)			(0.0057)	(0.0055)			(0.0055)	(0.0055)
Change in Real GDP				0.0002				0.0001
(percentage points)				(0.0005)				(0.0005)
Constant	0.0090***	-0.004	-0.0049	-0.0035	0.0090***	-0.0038	-0.0047	-0.0033
	(0.0010)	(0.0034)	(0.0037)	(0.0031)	(0.0010)	(0.0034)	(0.0036)	(0.0030)
Observations	213	205	196	187	213	205	196	187
Adjusted R^2	0.0855	0.1693	0.2069	0.2231	0.0814	0.1672	0.2113	0.2210

Notes: This table reports the results of two alternative sets of jurisdiction-level regression specifications, one without the citizenship-for-sale covariate, and one adding a Tax haven X CFS interaction covariate. Removing CFS as a covariate does not materially affect the estimates on other covariates. Adding the haven interaction reveals that the CFS effect is driven by the non-haven CFS jurisdictions, like Bulgaria and Serbia.

Figure 1.19: Count of jurisdictions signing FATCA IGAs, by year and model type

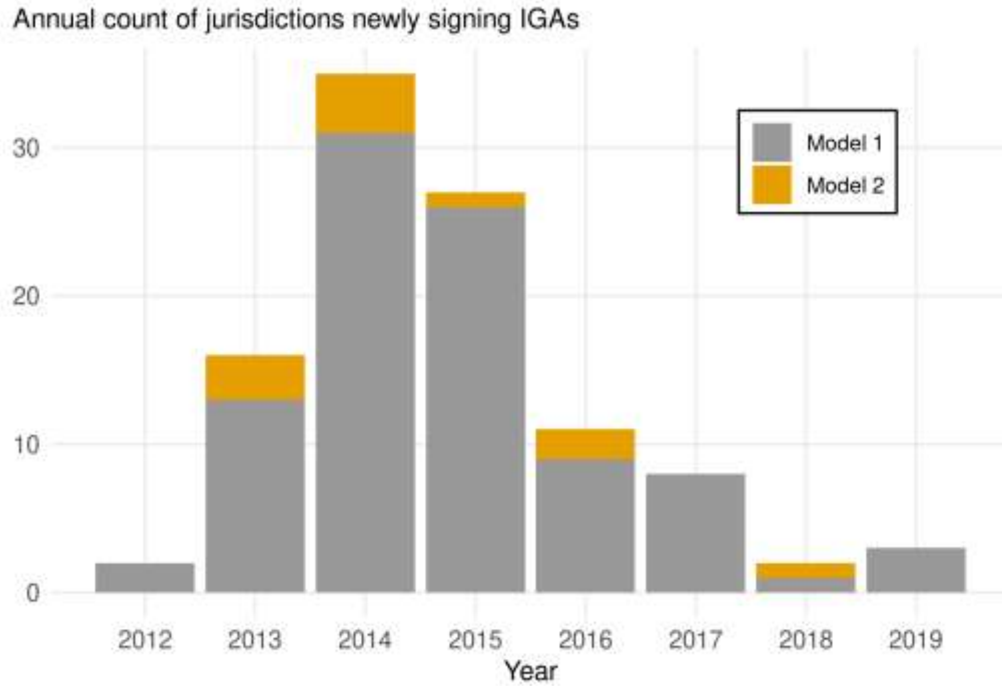

Notes: This figure plots the count of unique jurisdictions that signed a FATCA-related IGA with the U.S. in each year, separately for the two types of IGAs.

Figure 1.20: Jurisdictions by FATCA IGA status

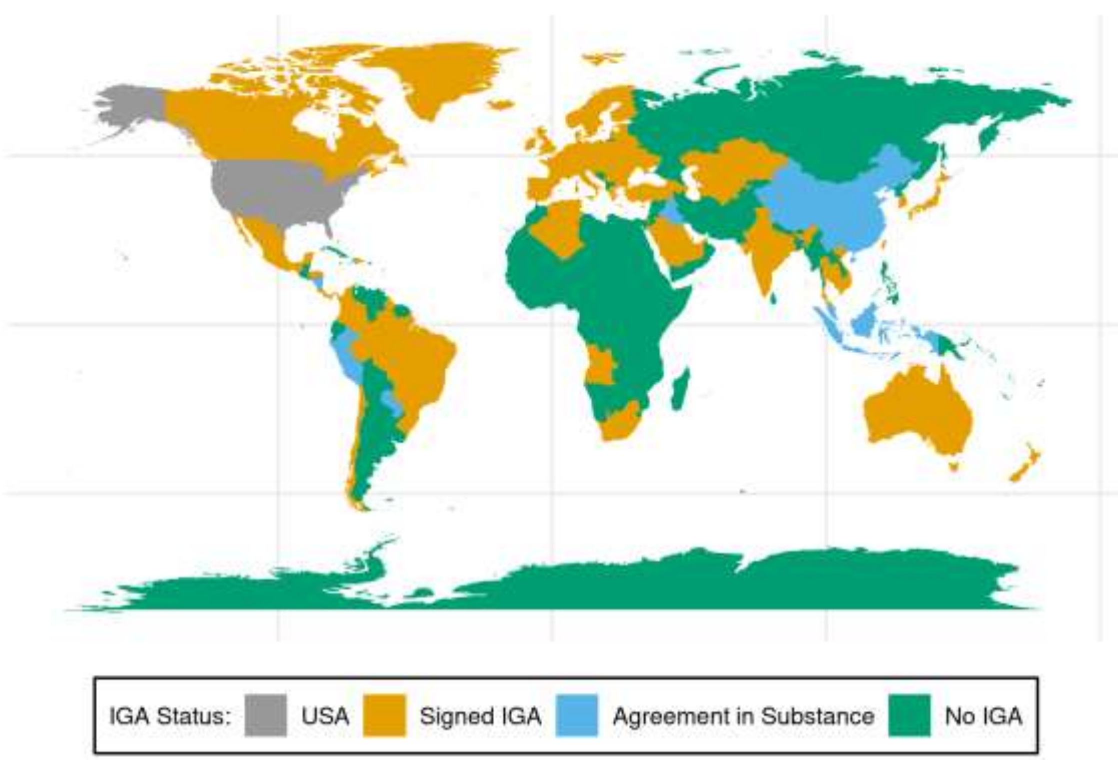

Notes: This figure shows world jurisdictions, shaded to indicate their FATCA IGA status, according to the U.S. Treasury. Jurisdictions with a signed IGA are shaded in orange; those listed as having an "Agreement in Substance" but not yet signed are shaded in blue. Those with neither a signed IGA nor an Agreement in Substance are shaded in green.

Figure 1.21: Annual share renouncing, alternate specifications

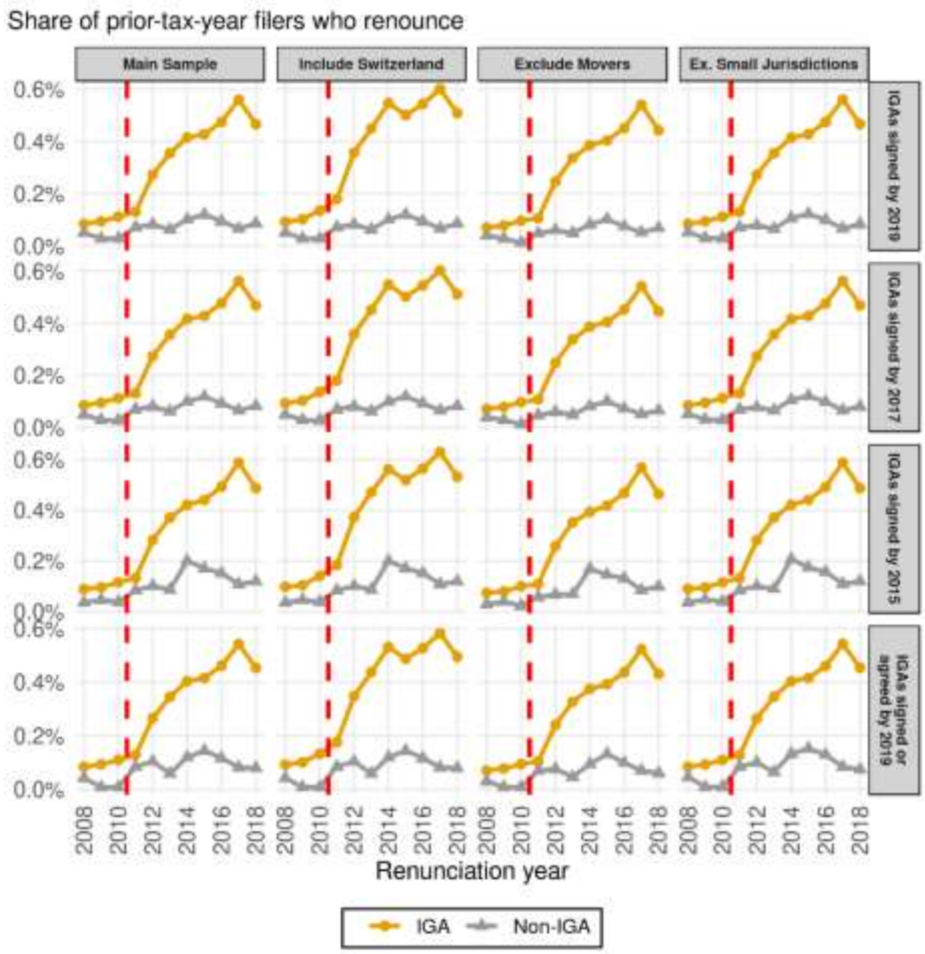

Notes: This figure plots the number of individual renouncing each year as a share of prior-year U.S. tax filers, from either IGA or non-IGA jurisdictions. The four rows correspond to four different definitions of "IGA jurisdictions": including those signed through 2019, 2017, or 2015, or adding in those with Agreements in Substance. The four columns correspond to sample definitions: the second column includes observations from Switzerland; the third column excludes Movers; and the fourth column restricts to jurisdictions with at least 100 U.S. tax filers in one year.

Table 1.15: Difference-in-difference results, alternate data specifications

Dependent variable:	Renounce in following year (100/0)			
	[1]	[2]	[3]	[4]
Sample:	Main	Include Switzerland	Exclude Movers	Exclude Small Jurisdictions
Post	0.0455*	0.0450*	0.0376*	0.0450*
	(0.0251)	(0.0250)	(0.0211)	(0.0265)
IGA jurisdiction	0.0465	0.0618*	0.0414	0.0449
	(0.0338)	(0.0367)	(0.0299)	(0.0340)
Post X IGA	**0.2579****	**0.3168*****	**0.2581*****	**0.2582****
	(0.1011)	(0.1023)	(0.0996)	(0.1016)
Individual covariates	Yes	Yes	Yes	Yes
Observations	4,690,000	4,835,000	4,686,000	4,664,000
Adjusted R^2	0.0013	0.0013	0.0012	0.0013
Mean dep. var.	0.312	0.367	0.291	0.313

Notes: *p<0.1; **p<0.05; ***p<0.01. The coefficient estimates on the Post, IGA, and PostXIGA covariates are shown under alternate sample definitions. Standard errors clustered by year and by jurisdiction are shown in parentheses. The four columns correspond to sample definitions: the second column includes observations from Switzerland; the third column excludes Movers; and the fourth column restricts to jurisdictions with at least 100 U.S. tax filers in one year.

Table 1.16: Difference-in-difference results, alternate IGA definitions

Dependent variable:	Renounce in following year (100/0)			
	[1]	[2]	[3]	[4]
Including IGAs signed through:	*2019*	*2017*	*2015*	*2019 and Agreements in Substance*
Post	0.0455*	0.0450*	0.0904***	0.0794***
	(0.0251)	(0.0242)	(0.0272)	(0.0201)
IGA jurisdiction	0.0465	0.0476	0.0555*	0.0727***
	(0.0338)	(0.0336)	(0.0311)	(0.0280)
Post X IGA	**0.2579****	**0.2591*****	**0.2235****	**0.2142****
	(0.1011)	**(0.1005)**	**(0.1006)**	**(0.0969)**
Individual covariates	Yes	Yes	Yes	Yes
Observations	4,690,000	4,690,000	4,690,000	4,690,000
Adjusted R^2	0.0013	0.0013	0.0014	0.0012
Mean dep. var.	0.312	0.312	0.312	0.312

Notes: *p<0.1; **p<0.05; ***p<0.01. The coefficient estimates on the Post, IGA, and PostXIGA covariates are shown under alternate IGA definitions. Standard errors clustered by year and by jurisdiction are shown in parentheses. The four columns correspond to four different definitions of "IGA jurisdictions": including those signed through 2019, 2017, or 2015, or adding in those with Agreements in Substance.

Table 1.17: Revenue-relevant summary statistics

		Split by prior U.S.-based filing		Split by time period	
	All	Mover	Dropper	1998-2010	2011-2018
Total count of renouncers	39,300	4,600	20,000	9,100	30,200
Count with TIN/SSN	27,900	4,600	20,000	6,500	21,400
Count with Form 8854	20,400	3,200	14,500	5,000	15,400
Tax liability one year prior to renunciation					
Count with non-zero liability	8,700	2,200	6,100	2,400	6,200
Among those with TIN: share with non-zero liability	31%	48%	31%	37%	29%
Among those with non-zero liability:					
Median liability ($)	$4,200	$11,800	$3,000	$9,300	$3,200
Mean liability ($)	$81,000	$183,500	$46,900	$137,500	$60,400
Total liability ($M)	$704	$404	$286	$330	$374
Median age at expatriation	52	51	52	50	52
Average annual tax liability over five years prior to expatriation					
Count with at least one year of non-zero liability	14,800	3,600	10,600	3,800	11,000
Among those with TIN: share with non-zero liability	53%	78%	53%	58%	51%
Among those with non-zero liability:					
Median liability ($)	$2,200	$8,000	$1,400	$6,200	$1,500
Mean liability ($)	$41,500	$94,000	$24,900	$71,700	$31,000
Total liability ($M)	$614	$338	$264	$272	$341
Median age at expatriation	52	50	52	49	53

Notes: pre-FATCA includes expatriations from 1998-2010; post-FATCA from 2011-2018. Values are rounded to the nearest 100 for disclosure purposes. Tax liability is net of the Foreign Tax Credit.

Figure 1.22: Annual count of renunciations, split by pre-renunciation U.S. tax liability

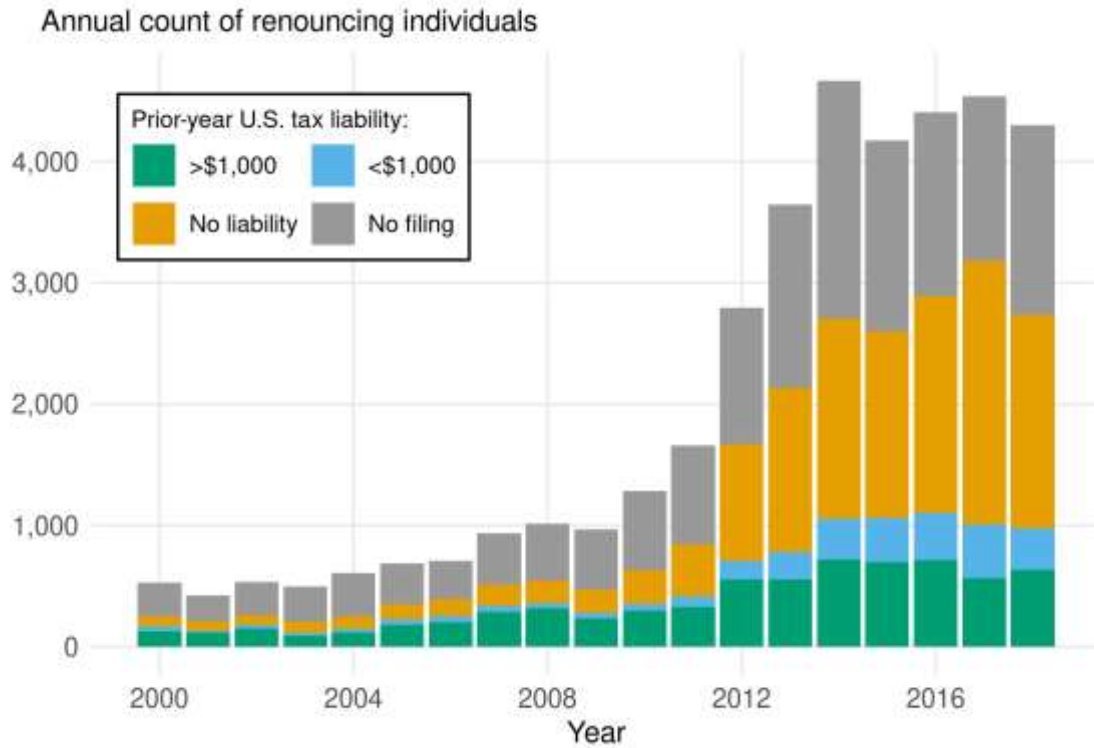

Notes: This figure plots the count of individuals renouncing each year, split by their U.S. tax liability in the year prior to renunciation. In gray are those without linked U.S. tax filings (mostly due to lack of TIN/SSN for linking). In orange are those with zero net liability; in blue those with <$1,000 in net liability; and in green, those with >$1,000 in net liability. Most of the recent increase in renunciation is from those without filings or liability.

Figure 1.23: Histogram of reported net worth around $2 million (extended range)

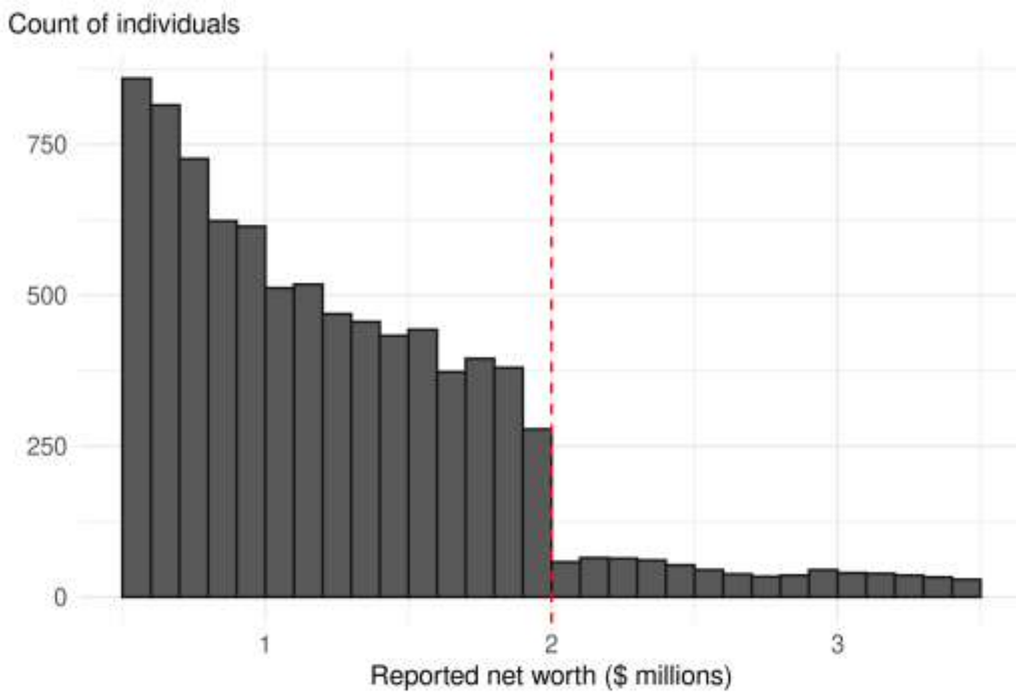

Notes: This figure plots the count of expatriates with past-5-years average income tax liability in each $10K bucket of liability relative to the covered expatriate threshold in their year of renunciation. The thresholds are shown above in Figure 1.25.

Figure 1.24: Histogram of average income tax liability relative to threshold

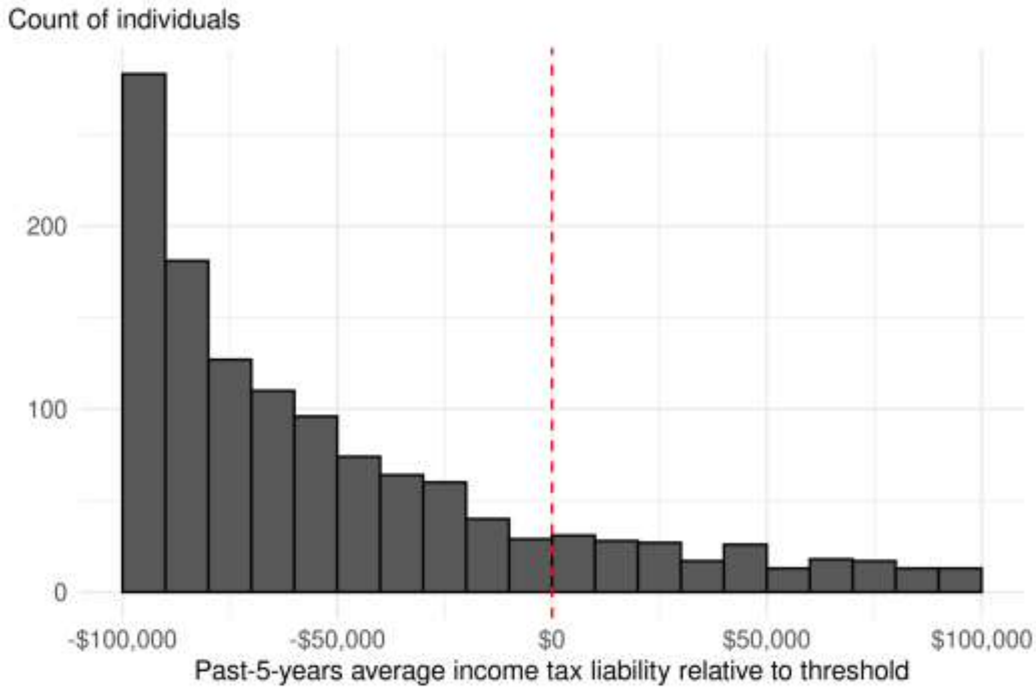

Notes: This figure plots the count of expatriates with past-5-years average income tax liability in each $10K bucket of liability relative to the covered expatriate threshold in their year of renunciation. The thresholds are shown in Figure 1.25.

1.9. Appendix B: Additional notes on U.S. citizenship renunciation

1.9.1. The process of U.S. citizenship renunciation

Section 349(a) of the Immigration and Nationality Act outlines the seven acts by which a U.S. national can voluntarily relinquish U.S. nationality: (1) obtaining naturalization in a foreign state once 18 years or older; (2) declaring allegiance to a foreign state once 18 years or older; (3) serving in the armed forces of a foreign state, either as an officer or engaged in hostilities against the U.S.; (4) serving a foreign government if that service requires foreign nationality or allegiance; (5) making a formal renunciation of nationality before a diplomatic or consular officer of the U.S. in a foreign state; (6) making a formal written renunciation of nationality in the U.S. (when the U.S. is at war and the renunciation is approved by the Attorney General); and (7) committing treason against or attempting to overthrow the U.S. government (8 U.S.C. §1481).

The fifth option is the main approach taken by those choosing to lose their U.S. citizenship. The required steps include preparing the necessary forms, meeting with a diplomatic or consular officer in a foreign state, swearing an oath of renunciation, and paying the renunciation fee. In most cases, renunciation requires two separate appointments at a foreign embassy or consulate. In the first appointment, the U.S. citizen is interviewed to confirm that renunciation is being done out of free will and not under duress. At the second appointment, an oath of renunciation is sworn. The current fee for citizenship renunciation is $2,350 (until mid-2014 the fee was $450, and there was no fee before 2010). Because of the recent increase in renunciations, some embassies and consulates have experienced backlogs of renunciation appointments, leading to delays or prompting some individuals to travel to other cities and countries to seek earlier appointments (Richards 2016). Various third-party firms offer services to U.S. citizens considering citizenship renunciation, promising to assist with the process and often targeting their marketing at high-wealth individuals.[44]

After completing these steps, the State Department processes the renunciation and sends the individual a Certificate of Loss of Nationality confirming the renunciation of U.S. citizenship. Those renouncing citizenship must also file a U.S. income tax form for the year in which they renounced citizenship and include Form 8854 to complete renunciation for tax purposes (and remit or make arrangements to remit any associated tax liability). Because expatriation is an individual process, each individual must file a separate Form 8854, even if filing Form 1040 with married filing jointly status.

Figure 1.25 below plots the changes over time in the income tax and net worth thresholds for covered expatriate designation, as well as the changes in the capital gains exemption available to those who are deemed covered expatriates and subject to the mark-to-market exit tax (since 2008).

[44] See, e.g., the Nomad Capitalist (https://nomadcapitalist.com/) or 1040Abroad (https://1040abroad.com/about/).

Figure 1.25: Statutory covered expatriate thresholds and gains exemptions over time

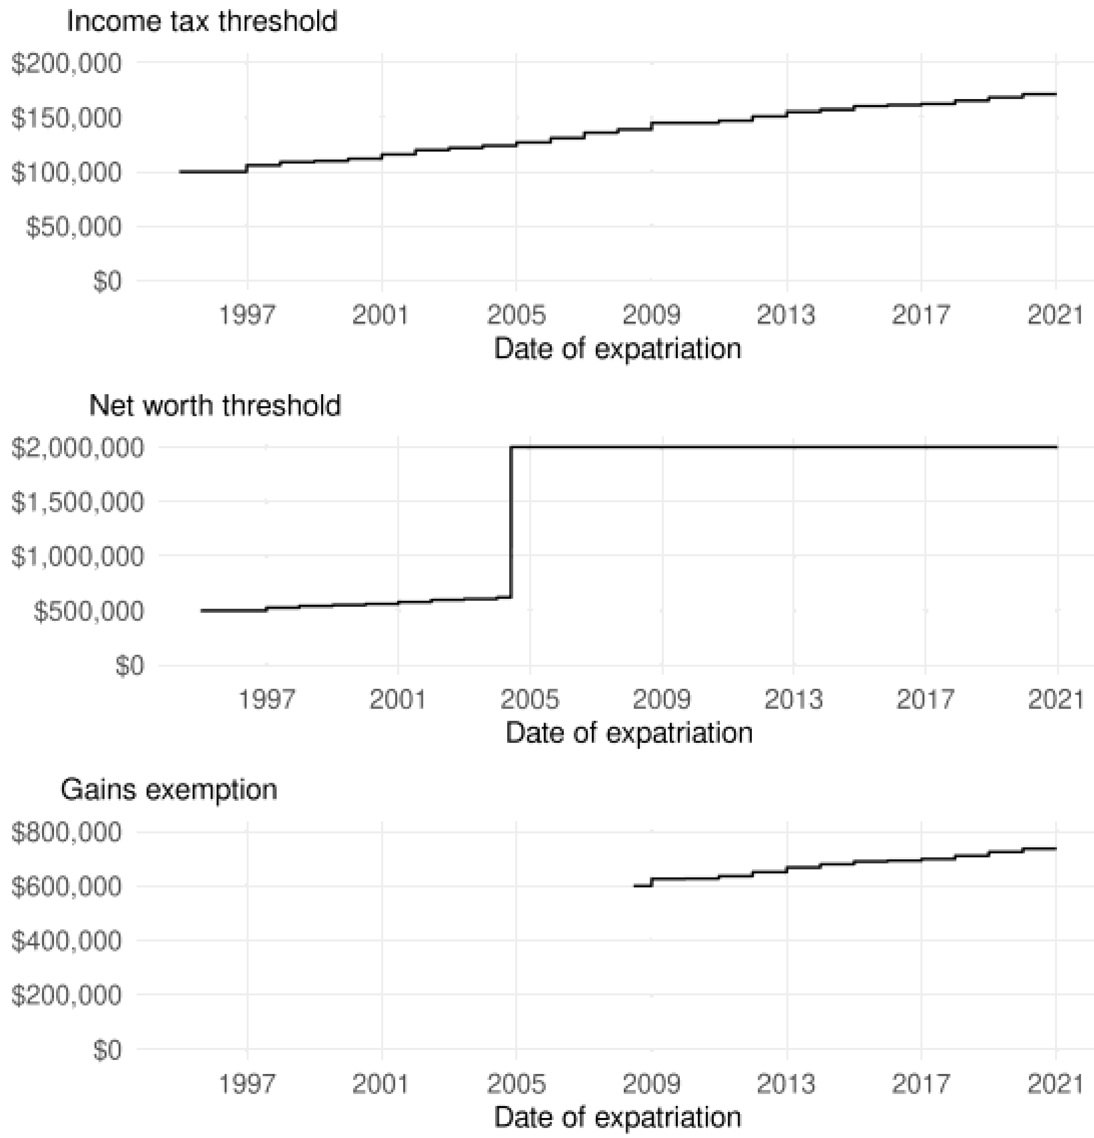

1.9.2. Prior U.S. history with citizenship renunciation

In the main text of this paper I focus on renunciations from 1998-2018, the years for which the IRS database with information on those renouncing U.S. citizenship is available and complete. Using outside sources, it is also possible to provide some context on renunciations covering a longer time period.

Annual counts of U.S. citizen renunciations are available for the years 1962-1994 from the State Department, as listed in a report discussing proposals for changes to the tax treatment of expatriation (Joint Committee on Taxation 1995). These can be paired with annual counts of individuals reported in the Federal Register as having relinquished citizenship, which are available for the years 1998-2020. Note that because of slight differences in the way numbers were tracked from year to year and the precise criteria for inclusion, these sources may not be exactly comparable with each other, nor with the counts I present above.[45] Still, all three capture a similar idea and allow for consideration of trends over time.

Figure 1.26 below shows the JCT and Federal Register series, with the JCT numbers in blue and the Federal Register numbers in green. The longer-term trend shows that renunciations were actually somewhat more common in the 1960s and 1970s, and had fallen to a relative low by the 2000s, before increasing in the past decade, as discussed above.

[45] For instance, the JCT report notes at p. 7 that there may be discrepancies between the definitions used for the yeas 1962-1979 and 1980-1994.

Figure 1.26: Annual count of renunciations (JCT and Federal Register)

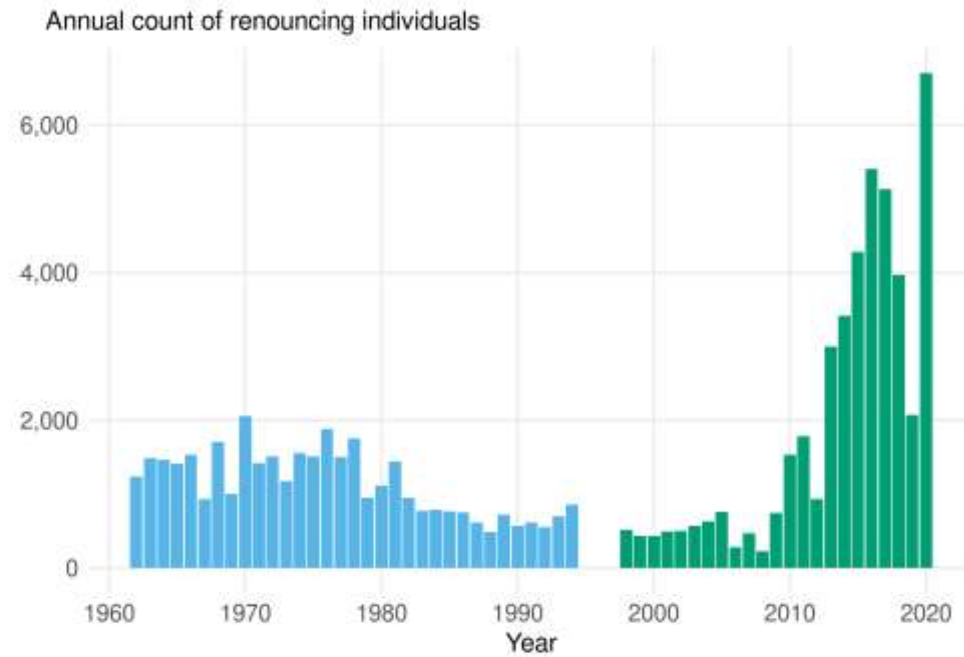

Notes: This figure plots the count of U.S. citizenship renunciations from two sources. For the years 1962-1994, the values are as reported in Joint Committee on Taxation (1995). For the years 1998-2020, the values are the count of names published in the Federal Register as the "Quarterly Publication of Individuals Who Have Chosen to Expatriate", required under IRC §6039G.

Chapter 2. Your Passport, Please: Incentive Effects of the IRS' Passport Certification and Revocation Process (with Alex Ruda, Joel Slemrod, and Alex Turk)[46]

2.1. Introduction

Collecting taxes that are owed but unpaid is a difficult issue facing tax agencies around the world. Traditionally, authorities have relied mainly on financial penalties to induce delinquent taxpayers to resolve their debts and, on some occasions, seizure of real assets or criminal prosecution. However, there is a third tool that is increasingly used, known as collateral sanctions. Blank (2014) defines these as measures that are applied in addition to formal tax penalties, rescind government-provided benefits or privileges, and are usually enforced by an agency other than the tax agency. Examples of collateral sanctions include restricting access to drivers' or professional licenses, and publishing the names or information of tax delinquents. Although empirical analyses of the effectiveness of monetary enforcement policies have recently proliferated (Slemrod 2019), little attention has been paid to evaluating the effectiveness of collateral sanctions.

[46] We thank Ricardo Perez-Truglia and three referees for extremely helpful and constructive comments. An earlier version of the paper was circulated with the title "Do Collateral Sanctions Work?" We thank Nick Gebbia, Tatiana Homonoff, Jeff Hoopes, Jan Millard, Stacy Orlett, and participants at the 2020 National Tax Association annual conference, 2021 IRS/TPC Joint Research Conference on Tax Administration, 2021 IIPF annual congress, and University of Michigan seminars for helpful comments. All data work for this project involving confidential taxpayer information was done on IRS computers by IRS employees, and at no time was confidential taxpayer data ever outside of the IRS computing environment. Organ is a student volunteer with the IRS through the Joint Statistical Research Program. Slemrod is an IRS employee under an agreement made possible by the Intragovernmental Personnel Act of 1970 (5 U.S.C. 3371-3376). The views and opinions presented in this paper reflect those of the authors. They do not necessarily reflect the views or the official position of the Internal Revenue Service. All results have been reviewed to ensure that no confidential information is disclosed.

This paper addresses that gap by examining a recent U.S. initiative restricting passport access for taxpayers with substantial tax debt. The primary sanction is denying an application or renewal, but can include revocation of an existing passport. We provide experimental and quasi-experimental evidence of the incentive effects of this program, including its effects on payments, the likelihood of entering into an installment payment agreement, and several other actions taxpayers can take to resolve the debt. We find small but positive, statistically significant effects on taxpayer compliance of the certification notice sent to all eligible taxpayers, and immediate and strong positive effects for a subset of certified taxpayers who make requests and are denied.

The collateral sanction we examine was enacted as part of the Fixing America's Surface Transportation (FAST) Act of 2015.[47] This Act required the Internal Revenue Service (IRS) to notify the State Department of taxpayers owing a seriously delinquent tax debt (initially $50,000 or more, now $54,000 or more as of 2021). The State Department must then deny passport applications or renewals for these taxpayers, and may also revoke existing passports. In February 2018, the IRS began notifying taxpayers of their certification and sending certifications of unpaid tax debt to the State Department. TIGTA (2019) discusses some issues with its implementation.

The analyses by the U.S. Government Accountability Office (GAO) and the Treasury Inspector General for Tax Administration (TIGTA) help put into perspective the enforcement potential of the passport sanction program. According to GAO (2011), in fiscal year 2008, of the 16 million passports issued, 224,000 were issued to individuals who owed a total of $5.8 billion in unpaid federal taxes. The TIGTA (2019) report provides a snapshot of the passport certification

[47] The idea was first introduced in fall 2011 as part of Sen. Barbara Boxer (D-CA)'s bill called MAP-21, a precursor to the FAST Act (Urist 2012). There is some precedent for restricting access to passports for certain individuals; the Conference Report on the FAST Act notes that under then-present law, the Secretary of State could "refuse to issue or renew a passport if the applicant owes child support in excess of $2,500 or owes certain types of Federal debts." (U.S. House of Representatives 2015).

program as of May 17, 2019. By that date, the IRS had certified nearly 400,000 taxpayers as having "seriously delinquent tax debt" under the FAST Act.[48] A total of $961 million in payments had been made to these taxpayer accounts, including $551 million in fully paid balances. As a result of these collections efforts, over 40,000 taxpayers had received decertifications for having paid in full or otherwise moving towards payment of the debt. This accounted for 40% of all decertifications through May 17, 2019, with the remaining 60% resulting from taxpayers' administrative actions falling into one of several categories, the most prevalent being temporary relief due to disaster declaration (27%), tax debts becoming uncollectible due to statute expiration (12%), or taxpayers qualifying for hardship exemptions (9%).

What these numbers do not reveal is the counterfactual—how much of this debt would have been paid off in the absence of the passport sanctions program—and therefore they do not provide an estimate of the causal impact of the passport program. In this paper, we take advantage of administrative tax microdata and a randomized controlled trial (RCT) embedded in the implementation of the program to provide credible causal estimates of its impact. During the program rollout, 5% of taxpayers eligible for passport certification were randomly selected based on Social Security numbers and temporarily held out as a control group. We estimate the direct causal effect of passport certification on the full sample of about 265,000 taxpayers who were eligible for certification during the first year of the program.

Our results show that passport certification causes a non-trivial fraction of the certified taxpayers to take action leading to decertification. Using an instrumental-variables approach to address treatment migration in the RCT, we find that certification leads to a two-percentage-point increase in the probability of taking a compliance action that would lead to decertification (on a

[48] The report notes that this includes some repeat certifications, where a taxpayer was certified, de-certified, and later certified again. In this paper, we focus on first-time certifications.

base of 17%). The most common such action is entering into an Installment Agreement (IA), in which the taxpayer commits to paying down their tax debt over time. The effect of certification on starting a new IA is 1.3 percentage points (on a base of 7.6%). These effects are especially striking given that less than half of Americans hold passports, and only a fraction of those that do would need to renew or make changes within a few years. Although we do not observe passport holding directly, we can observe characteristics that correlate with it, and find that the effect of certification is higher for taxpayers with higher income and for those with evidence of some foreign activity.

We also study the response of approximately 10,000 certified taxpayers who made passport-related requests after being certified (i.e., notified that their request might be denied) and were denied**.** This is certainly a very special group, who may not have received or read the certification notices, or may have received the notifications but guessed that the IRS was bluffing. Nonetheless, we find clear evidence that a substantial fraction of those having a passport request denied immediately take action to become decertified and therefore once again eligible for passport access. These actions also lead to additional payments; those with requests denied paid $9,000 more, on average, over the following year, compared to a randomly selected control group of certified taxpayers without denied requests.

We conclude by using our findings to assess the revenue and welfare implications of the passport program. The results of the RCT and State Denied analyses show that the passport program has significant incentive effects and that these lead to substantial additional revenue, far exceeding the administrative costs of running the program. However, it is also important to consider the private costs of the program; although more difficult to quantify, the disutility of certified taxpayers who are unable to apply for or renew their passports should be considered when evaluating the program.

Our paper contributes to a growing literature studying the effects of non-financial policies on tax compliance. Empirical work has analyzed the effects on tax compliance of shaming tax delinquents by publishing their personal information (Perez-Truglia and Troiano 2018, Dwenger and Treber 2019, Angaretis, et al. 2021); and professional license suspensions (Kenchington and White 2021). Theoretical work has shown that collateral sanctions could have incentive effects that are more salient than monetary penalties and may affect reputation (Blank 2014), and could allow for more targeted enforcement by affecting consumption opportunities that are correlated with earnings potential (Kuchumova 2018, Kuchumova 2021). We contribute to this literature by studying a notable and, to our minds, fascinating collateral sanction—restricting access to passports—about which little is known. Our analysis also strengthens the connection between debt enforcement policy and the much larger literature on tax evasion enforcement. The policy we study is designed to accelerate the collection of tax debt, while the great majority of empirical analyses of tax enforcement actions address their impact on tax evasion. These are not the same thing, in part because tax debt can accumulate even if no evasion has occurred, by the nonpayment of agreed-upon tax liability. But they are conceptually related in an important way. Taking into account whether tax debts are collected adds nuance to the classic Allingham and Sandmo (1972) model of tax evasion.[49] This model implicitly assumes that all agreed-upon debts are collected. If that is not true, then changes in the debt collection efficacy affect both the expected present value of the tax remitted per dollar of reported income, and the present value of the penalties remitted per dollar of unreported income uncovered in an audit. Thus, the perception of debt collection efficiency should affect the magnitude of tax evasion.

[49] Hallsworth et al. (2017) note the connection between the Allingham-Sandmo model and the enforcement of tax debt.

The rest of the paper proceeds as follows. Section 2.2 discusses previous research on the use of collateral sanctions and other tax compliance tools. Section 2.3 provides institutional details about the passport program, and Section 2.4 describes the data we use. Section 2.5 describes our RCT analysis of the direct effects of certification. Section 2.6 considers the effect of denied passport-related requests for the subset of certified taxpayers who made such requests. Section 2.7 discusses the policy implications of our results, and Section 2.8 concludes. The Appendices include additional figures and tables (Section 2.9, Appendix A); comments on potential indirect effects of the passport program (Section 2.10, Appendix B); and sample certification and decertification notices (Section 2.11, Appendix C).

2.2. Previous literature

Until the passport provisions of the FAST Act, federal use of collateral sanctions for tax purposes was limited—failure to pay taxes may result in the loss of ability to apply for FHA mortgages and enter into contracts with the federal government. States have been more active users of collateral sanctions. Some suspend drivers' licenses and vehicle registrations (e.g., California), revoke law and professional licenses (e.g., Minnesota and Wisconsin), or deny hunting and gaming permits (e.g., Louisiana). Criminal convictions result in a wide range of collateral sanctions.

Blank (2014) offers several reasons why a collateral sanction may be more effective than a monetary penalty: it may be more salient; as a denial of an existing service, it may trigger loss aversion; and it may affect one's reputation. Kuchumova (2018) formalizes a rationale for their use: by affecting consumption and providing enforcement targeted to a group, collateral tax sanctions can allow the government to impose punishment correlated with an individuals' earning potential. On the other hand, as Polinsky and Shavell (2000) note, for non-monetary sanctions such as a collateral sanction (e.g., imprisonment), the private cost to the sanctioned party is not matched by an equal transfer to the government as it would be for a monetary penalty. In our setting, certified taxpayers experience the disutility of restricted passport access but the government does not benefit from this disutility. From the social planner's perspective, this cost of collateral sanctions may be partially offset if the resource cost of administering the public good affected by sanctions is reduced by less demand due to the sanctions (in our case, administering passports may be less costly if demand for passports is reduced by tax debt-related certifications and revocations). In addition, the authority to impose these collateral sanctions could reduce the administrative costs of debt collection if they are substitutes for more resource-intensive enforcement actions.

The passport authority is not the only policy lever the IRS can use to collect tax debt. A "Notice of Federal Tax Lien" (NFTL) can be filed for tax assessments that are not paid. The filing makes the otherwise private debt public information, and helps to establish the government's lien on the taxpayer's assets. This can impact the individual or business taxpayer's credit rating and their ability to sell or refinance assets. Turk et al. (2016) use an event study to examine the impact of NFTL policy changes that were put in place during the great recession. Collins et al. (2018) study the NFTL impact versus direct contact in a randomized control trial. Most of the previous research points to significant effects of NFTL filing policy on debt resolution and (to a lesser degree) indirect effects on payment compliance.

As noted above, while there is some theoretical literature on collateral sanctions, there is relatively little empirical work quantifying the effects of non-financial tax debt enforcement tools in practice. One tool that has received attention is publicizing the names and personal information of tax debtors. Perez-Truglia and Troiano (2018) conducted a letter RCT, randomizing whether letters were sent, as well as the content of sent letters, to tax debtors whose information had already been published by state tax agencies. Letters increased the probability that low-balance tax debtors (those with debts below $2,500) would leave the list, but had no effect on high-balance tax debtors**.** Dwenger and Treber (2019) study the introduction of a delinquent taxpayer public disclosure program for corporations and self-employed individuals in Slovenia. They find that the threat of publication leads both corporations and the self-employed to reduce tax debt. Publication itself led to further reductions in tax debt, though of a smaller size**.** Angaretis et al. (2021) study California's "Top 500" tax delinquent publication scheme, which publicizes the names of the 500 taxpayers with the largest unpaid state tax liabilities. They find that about a quarter of taxpayers resolve their debts when notified that they will shortly be published, suggesting publication is perceived to be costly, at least for these high-liability taxpayers. Taken together, these studies show that shaming can be an effective tool for collecting unpaid tax liabilities. There is also some evidence on the use of license suspensions as a tax enforcement tool, which is even more directly related to the passport program we study here. Kenchington and White (2021) study Missouri's policy of suspending the professional licenses of those who are either delinquent in remitting state income taxes owed or have not filed for three years, finding that license suspensions are relatively common for lower-income professions. They interpret this as evidence that, for some, tax noncompliance is driven by financial constraints rather than unwillingness to pay. Angaretis et al. (2021) also provide some evidence of the effect of license suspension because, in addition to publicizing the names of top tax debtors, California's "Top 500" program can also result in suspension of professional, driver's, and other licenses. Comparing the payment trends of published taxpayers with and without licenses, the authors conclude that the threat of license suspension induces a positive compliance response.

Our paper adds to this growing literature on the incentive effects of non-monetary policies by providing experimental and quasi-experimental evidence of the effects of the IRS' passport certification and revocation process.

2.3. Institutional background

2.3.1. Certification and notification

The first step in the passport program is a determination of which taxpayers are "certified" to have eligible debt over the threshold. Certification applies to "modules" – basically, a taxpayer-tax year combination. If the sum of tax due and assessed interest and penalties on all eligible modules for a given taxpayer exceeds the threshold, all of the eligible modules are certified. For a module to become eligible, it must constitute "seriously delinquent tax debt" by meeting at least one of two inclusion requirements, and not meet a set of exclusion requirements. The two inclusion requirements require that the module either has had a Notice of Federal Tax Lien filed and that the associated Collection Due Process (CDP) hearing rights have expired, or that a Notice of Levy has been issued. For either of these inclusion requirements to be met means that the taxpayer has been through many months, and often years, of the traditional enforcement process, and has been sent numerous notices encouraging resolution of their outstanding debt. The exclusion requirements exempt modules that meet any of a set of detailed criteria, the most common of which are being part of an approved or pending Installment Agreement or Offer-in-Compromise, about which we say more below.[50]

The debt threshold was set at $50,000 initially and is indexed for inflation, rising to $51,000 in 2018, $52,000 in 2019, $53,000 in 2020, and $54,000 in 2021. The relevant balance for this

[50] For more details on specific eligibility and ineligibility criteria, and the treatment of certain special cases, see Internal Revenue Materials section 5.19.25 (https://www.irs.gov/irm/part5/irm_05-019-025).

threshold is unpaid *assessed* balance (tax, interest, and penalties).[51] Interest and penalties that accrue over the time after assessment is made do not influence whether a taxpayer is eligible for certification. Once modules have been certified, a taxpayer must pay off or otherwise deal with *all* of the certified module debt in order to be decertified. For example, suppose a taxpayer has eligible modules from three tax years with assessed amounts of $25,000, $20,000 and $10,000, for a total of $55,000. This total breaches the threshold and each of the three underlying modules is certified. The taxpayer must then resolve each of these three to become decertified; for example, only resolving the $10,000 module (resulting in $45,000 of remaining eligible balance, below the threshold) will not suffice.

About two weeks after certification, both the IRS and the State Department send letters to the taxpayer notifying them of their certification.[52] The IRS letter is called a 508C Notice; it indicates the total balance identified as seriously delinquent, broken down into tax debt, penalties, and interest. The letter informs the taxpayer that as a result of certification:[53]

> If you apply for a passport or passport renewal, the State Department will deny your application and will not issue a passport to you or renew your current passport. If you currently have a valid passport, the State Department may revoke your passport or limit your ability to travel outside the United States.

[51] Taxes are typically assessed when the return is filed. For timely filed returns, interest and penalties accrue from the due date of the return for any unpaid amounts, but these accrued amounts are not included for passport program eligibility. For late filed returns (or returns with audit assessments) any interest and penalties due when the tax is assessed *are* included in the assessed balance, but subsequent accruals are not.

[52] The delay between certification and notification by the IRS allows the IRS and State Department to send their letters at roughly the same time.

[53] A sample 508C Notice is included in Appendix C. Beginning in 2016, similar language informing taxpayers about the passport program began to be included on other IRS notices sent to taxpayers as part of the normal collections process. Although this information was not the focus of those notices, its inclusion may have alerted some taxpayers about the potential for their future certification. To the extent this happened, our analysis of the effect of the passport program may represent a lower bound, as those individuals most responsive to the program may have seen references to it in earlier notices and responded preemptively.

Appendix Figure 2.4 shows the monthly count of certifications, separately counting first-time certifications (in gray) and subsequent certifications for the same taxpayer (in orange).[54] When the program began there were about 200,000 taxpayers who were eligible for certification, and most of these taxpayers were certified in waves over the months of March-July 2018. After this initial rollout, the program has maintained a regular certification of several thousand newly eligible taxpayers each month. Some cases involving "complex debt"—those with aggregate debt from multiple filing statuses—were not included in the initial rollout of the program due to the more complicated programming needed to ensure accurate certification. In early 2020, these complex debt cases were added to the program. Passport certification was paused in April 2020 due to COVID-19.

For most certified taxpayers, the consequence of certification would be the denial of their application for a new passport or for renewal of a passport. In certain cases, the IRS may recommend that the State Department revoke an existing passport; the State Department makes the final decision. As of January 2020, about 10,000 certified taxpayers had tried to do something passport-related and were denied (including new applications, renewals, and modifications). We study these taxpayers' subsequent behavior in Section 2.6.

2.3.2. Decertification

A certified taxpayer can get decertified by completely satisfying all of their certified tax debt, either by paying in full, by initiating an Installment Agreement (IA), or by successfully making an

[54] Subsequent certifications can happen for a number of reasons. For example, a taxpayer with existing modules that are certified may accrue new tax debt, which is subsequently certified. Alternatively, a taxpayer may take action to become decertified on previously certified debt, and then because of later actions become eligible again for certification.

Offer in Compromise (OIC).[55] Previously certified taxpayers can also subsequently become ineligible for inclusion in the program by newly meeting a statutory exclusion criterion, such as by being designated as Currently Not Collectible (CNC)-Hardship, or by filing for bankruptcy. Appendix Figure 2.5 shows the monthly count of decertifications.

How to get decertified is a question that many certified taxpayers themselves may have asked. As Appendix Figure 2.6 shows, strong patterns are visible in traffic to the IRS webpage describing the passport program.[56] The webpage was first published in January 2017, and several news articles linked to it appeared on February 3, 2017, resulting in the largest single day of page visits. Visits also increased significantly in June and July 2018, when the largest number of certifications and notifications occurred.

2.4. Data

The main source of data for this study is a set of IRS databases tracking unpaid tax assessments, taxpayer activity, and IRS notice issuances. For each module (a taxpayer-tax period combination), we observe monthly the balance due, including separately both assessed and accrued balances. We also observe, monthly or weekly, taxpayer activity concerning a number of outcomes of interest, including payments, entering into new Installment Agreements, proposing and having accepted Offers-in-Compromise, and designation as CNC-Hardship or having filed for bankruptcy. Finally, we observe records indicating which modules are certified and decertified, and when.

[55] Installment Agreements are plans to pay off a tax bill over an extended period; they must be approved by the IRS. Offers in Compromise allow a taxpayer to settle her tax debt for less than the full amount; they must also be approved by the IRS.
[56] https://www.irs.gov/businesses/small-businesses-self-employed/revocation-or-denial-of-passport-in-case-of-certain-unpaid-taxes

We do not know which taxpayers have passports, which of those are soon to expire, and which taxpayers do not have passports currently but plan to get them in the future. The State Department provides data on passport issuance by state, and from that data one can infer some patterns in the demographics of passport holders. We are able to match taxpayers to their tax filings, and use that matched information to proxy for passport ownership (for example, by noting that those with higher income are more likely to have or want passports).

Because not all taxpayers have passports, and among those who do only some will require changes or renewal within a few years, one should not expect the passport program to change the behavior of every certified taxpayer. Table 2.1 provides some context for the number of tax debtors and amount of tax debt that could be expected to be materially affected by the program, as of May 2018 (when the program was rolled out). If we assume that the fraction of individuals with passports in the country as a whole, 42%, applies to eligible debtors, and that about 20% of passport-holders seek renewal in a two-year period, then about 30,000 certified taxpayers would face a refused renewal application within that time period. Making the same assumptions about the dollar value of debt implies that about $6 billion of tax debt would be immediately impacted, or less than 3% of the total outstanding debt. Although these are of course just back-of-the-envelope approximations, they highlight that the passport program likely did not immediately impact all of the total tax debt in a substantial way. One caveat to this argument is that it assumes taxpayers understand that certification will probably only affect them if and when they apply for passport renewal; if, instead, taxpayers perceive certification to mean that their active passports will be revoked, then one might expect a stronger or quicker response.[57]

[57] Some news articles about the passport program may have led taxpayers to think that revocation was likely. See, for example, “IRS to Revoke 260,000 American Passports” (Gonçalves 2018) or “If You Owe Taxes to IRS, You Could Lose Your Passport and Your Ability to Fly” (Phillips Erb 2017).

Table 2.1: Estimate of potentially treated population size

	Taxpayers (#)	Debt ($B)
Individual income tax filings for 2017	150,690,787	
Individuals with outstanding tax debt in May 2018	16,811,000	252.28
Assessed debt > $50 thousand	768,000	157.92
Eligible for certification	348,000	68.32
Having a passport	146,160	28.69
Passport needing renewal in next two years	29,232	5.74

Notes: This exercise assumes that 42% of individuals have passports (based on the State Department's count of valid U.S. passports in circulation relative to the U.S. population in 2018), that 20% of passports require renewal in the next two years (most passports are valid for 10 years), and that passport holding and renewal are uniformly distributed. Debt includes assessed balance due, penalties, and interest (the amounts that affect passport program eligibility), but excludes accrued interest and penalties (the amounts that do not affect passport program eligibility).

2.5. Effects of certification

Certification is the main policy treatment under the passport program, so we begin our analysis by studying its effects. We take advantage of a randomized controlled trial (RCT) during the rollout of the passport program to estimate the direct causal effect of certification.[58] Below we begin by describing the population of eligible tax debtors and the details of the RCT, then show the effect of certification graphically, and follow that with regression analysis. We conclude this analysis by considering treatment heterogeneity along two key dimensions. First, we consider whether responses to certification differ between pre-existing cases (that would have been eligible in years past, had the program existed then) and new cases (that become newly eligible after the program rolled out). Second, we consider whether responses differ when splitting taxpayers along several characteristics that are apparently correlated with passport holding, including income, age, and geographic location.

[58] We focus in this paper on the direct effects of the passport program, that is the response of taxpayers once they are certified or have passport requests denied. The passport program could also have indirect effects, for instance by inducing taxpayers with tax debt below the eligible threshold to take actions so that they do not become eligible for certification. It could also affect the efficacy of other enforcement tools, like liens and levies, that are prerequisites for passport certification. We discuss these potential effects in Appendix B.

2.5.1. RCT analysis

Identifying the impact of certification is facilitated by the fact that it was implemented with an embedded RCT. During the first year of the passport program, Social Security numbers (SSNs) were used to randomize whether a certification-eligible taxpayer would actually be certified or not; 5% of the sample of taxpayers with eligible debt were “held out” from the certification process during the time period studied. The behavior of the holdout group compared to those taxpayers “treated” with a certification notification can provide evidence of the causal impact of certification, as the two groups are on average identical in all characteristics and thus would have on average behaved identically in the absence of the program.

One complication to this strategy arises, however. Due to certain actions or changes in their account characteristics, some of the control-group taxpayers were subsequently identified as being eligible for recertification (even though they were in the control group) and on these subsequent occasions were certified. In the terms of Angrist (2005), the RCT implemented here is subject to “treatment migration,” wherein some taxpayers in the control group obtain the treatment (certification). We discuss below how we address this issue.

We begin by presenting visual evidence of the effect of certification. Figure 2.1 shows the difference in certification rates over time between the treatment group (who were eligible for certification and by SSN were intended to be certified) and the control group (who were eligible for certification and by SSN were intended to be held out and have their certification delayed). By March 2019, 100% of the treatment group had been certified, while about 20% of the control group had been certified by that time, thus leaving about 80% of the control group who had not “migrated.”

Figure 2.1: RCT analysis, comparison only by SSN groups, certification

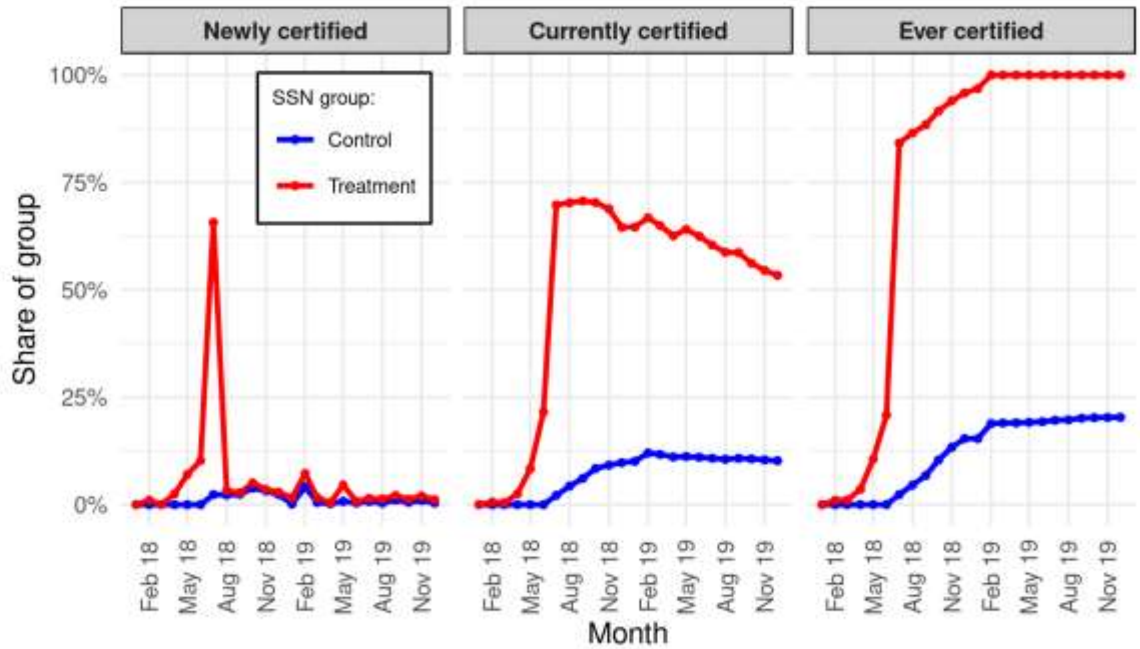

Notes: This figure compares taxpayers who by SSN were assigned to the treatment or control group during the rollout of the passport program. A wrinkle in the program implementation led roughly 20% of control group taxpayer to migrate into treatment (certification) over time.

Figure 2.2 shows some initial evidence that certification has a small positive effect on the probability of taxpayers taking positive compliance actions to resolve their debt. The figure compares the share of each SSN group that newly takes any compliance action leading to decertification. This includes starting a new Installment Agreement (IA); having an accepted Offer-in-Compromise (OIC); being newly designated as Currently Not Collectible (CNC) due to Hardship; filing for Bankruptcy; or having balance due go to zero, either by making full payment or having the balance abated. The jump in July 2018 is related to eligibility for certification; to be eligible, an individual must not have taken such compliance actions, and the majority of the RCT population was certified in July 2018. This is a “mechanical” reason to expect the probability to increase after July 2018. The important takeaway from this figure comes not from the level, but rather the comparison between the treatment (red) and control (blue). The figure shows a small but clearly higher rate of positive compliance actions among the treatment group, suggesting certification induces some individuals to resolve their tax debt.

Figure 2.2: RCT analysis, comparison by SSN groups, any compliance action

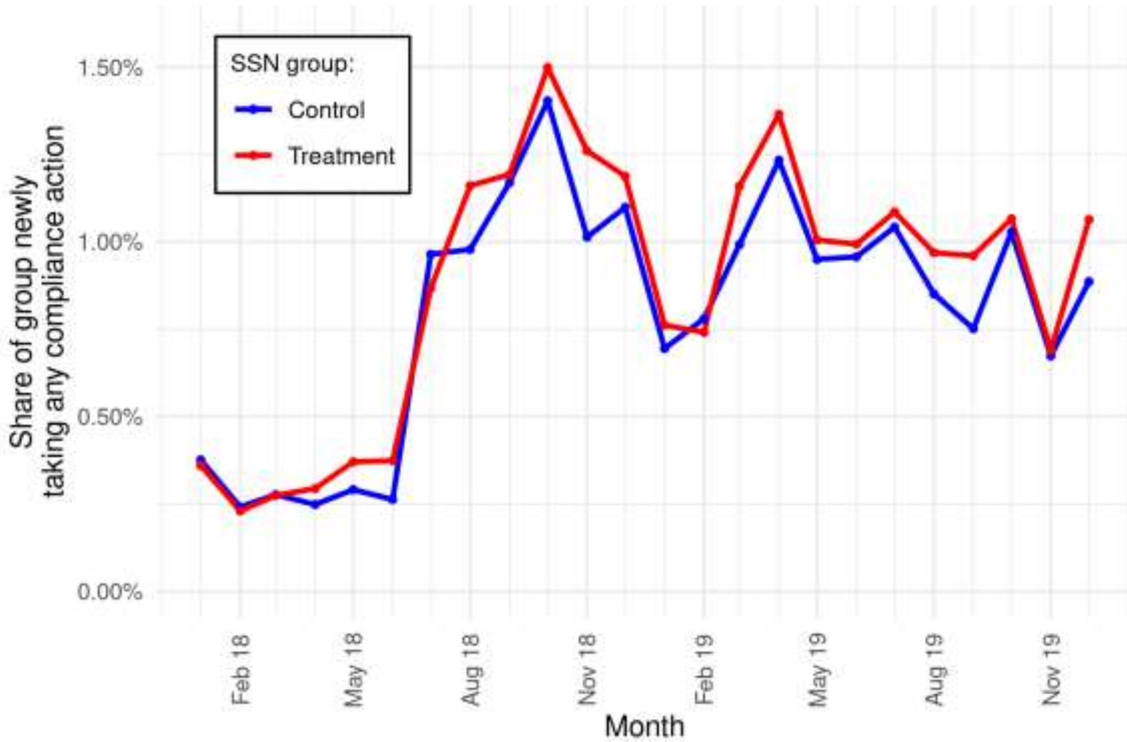

Notes: This figure compares the share newly taking a compliance action that will result in decertification for taxpayers who by SSN were assigned to treatment or control groups during the passport program rollout. All of the treatment group was certified, while approximately 20% of the control group migrated into treatment during the twelve-month RCT phase. The difference between the two groups reveals the effect of certification, with a small but clearly higher rate of actions taken among the treatment group. The compliance actions include starting a new Installment Agreement (IA); having an accepted Offer-in-Compromise (OIC); being newly designated as Currently Not Collectible (CNC) due to Hardship; filing for Bankruptcy; or having balance due go to zero, either by making full payment or having the balance abated.

We next consider separately the primary actions a taxpayer can take to get decertified. Appendix Figure 2.7 shows that there is a clear difference in the share newly taking action (top panel) and currently holding modules in the corresponding statuses (bottom panel) for IAs and CNCs. Certified taxpayers are more likely to take these actions, and to remain in these statuses over time. There appears to be a small difference in OICs, but no apparent difference in

Bankruptcies. Appendix Figure 2.8 shows that, in contrast, there is no apparent difference in the share of debtors newly resolving in full or the cumulative share that has fully resolved over time.

Because of the treatment migration, we cannot simply add up the estimated cumulative difference between what we have thus far labelled the treatment and control groups. To obtain estimates of the causal effect of certification, we follow Angrist (2005) and apply an instrumental-variables approach. The RCT-style implementation happened from the beginning of the program, in February, 2018, through February, 2019. We collapse the monthly variation in certification into a single dummy that equals 1 if a taxpayer was certified during this RCT phase, and 0 otherwise. The sample includes all those who are either certified or held out in the control group, during the RCT phase. For each of these taxpayers, we define the following:

- Y_i = outcome for individual *i* (e.g., change in balance)
- D_i = treatment dummy, 1 if individual *i* is certified during RCT phase, 0 otherwise
- Z_i = instrument dummy, 1 if individual *i* has a treatment SSN; 0 otherwise

We wish to estimate the following structural equation:

$$Y_i = X_i'\beta + \alpha D_i + \epsilon_i$$

To account for treatment migration, we begin by estimating the following first stage:

$$D_i = X_i'\pi_0 + \pi_1 Z_i + \eta_i$$

The reduced form, or second, stage is:

$$Y_i = X_i'\beta + \alpha[X_i'\pi_0 + \pi_1 Z_i + \eta_i] + \epsilon_i$$

$$= X_i'\delta_0 + \delta_1 Z_i + \nu_i,$$

allowing us to recover the causal effect of certification as follows:

$$\alpha = \delta_1/\pi_1.$$

This approach satisfies the assumptions necessary for a valid instrumental-variables method: within this population, (1) a taxpayer's SSN is correlated with certification and (2) the SSN only affects outcomes through its effect on certification.

Turning to the estimation itself, we first apply a few restrictions to the dataset.[59] To exclude outliers, we filter the dataset based on December 2017 characteristics (unaffected by the passport program rollout). Specifically, we exclude taxpayers with a total assessed balance due of more than $1 million, with a maximum module age (i.e., the age of their oldest tax-year-debt) more than 12 years, and those showing extreme numbers of modules.[60] Consistent with previous research on tax debt collection showing that debt resolution activity is not uniform across the population, we include a number of covariates based on December 2017 balances: total assessed balance, maximum module age, the share of the total balance older than 9 years[61], an indicator that the taxpayer had unfiled returns for at least one prior tax year, and the number of existing modules with balance due. We also include dummy variables for the source of assessment (telling us about how the debt balance started, e.g., voluntarily filed return, return secured from a non-filer, audit, etc.), status (telling us where the case is currently assigned in the collection process: currently assigned to field Revenue Officers or the Automated Collection System, inventories that are awaiting assignment, and inactive inventories), and taxpayer income type (three dummies indicating that the majority of income comes from wages, from sole proprietor income (Schedule C), or from interest, dividends, and capital gains; those with the majority of income from other sources are the excluded dummy). Table 2.6 in the Appendix presents summary statistics for the data used in the main RCT analysis.

[59] We test the sensitivity of our results to these restrictions and find little difference when adjusting them. Appendix Figure 2.9 shows the coefficient estimates under alternative data filters.

[60] We distinguish between Form 941 modules (Trust Fund Recovery Penalties assessed on a corporate officer for unpaid Form 941 liabilities, which are quarterly) and all other modules (which are annual). In our main specification, we exclude those taxpayers with more than 10 annual modules, or more than 40 quarterly modules.

[61] This allows us to control for the fact that the statute of limitations on most tax debt is 10 years. In some cases this means that a taxpayer's total balance due can fall due to the time limit being reached, and not because of action they have taken.

We test the effect of certification on a number of outcomes as of December 2019, thus allowing several months for behavior to be affected. In particular, we examine several individual actions that lead to decertification: IAs, OICs, CNCs, and Bankruptcies, as well as full resolution (i.e., having total assessed balance go to zero, by payment or by abatement). For each of these, we study the effect on taking the action by December 2019.[62]

[62] We test the sensitivity of the results to our choice of end month and find little difference when choosing earlier end months. Appendix Figure 2.10 shows the coefficient estimates for the range of end months from March 2019 to December 2019. We do not include outcomes in 2020 for two reasons. First, the rollout of enhanced programming for complex debt certification in early 2020 resulted in many taxpayers getting newly certified or recertified, which especially affects our estimates conditional on remaining uncertified. Second, and more importantly, the COVID-19 pandemic led to the pausing of certifications in April 2020, and likely affected taxpayer behavior as well.

Table 2.2: *RCT IV analysis, results for newly taking various decertification actions*

	Taking new action any time Mar '18-Dec '19				Fully resolved as of Dec '19		Any of six listed actions
	IA	OIC	CNC	Bankruptcy	By payment	By abatement	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Certified	**0.0131**	**0.0029**	**0.0049**	**-0.0004**	**0.0016**	**0.001**	**0.0205**
	(0.0026)	(0.0017)	(0.0020)	(0.0014)	(0.0013)	(0.0009)	(0.0037)
Covariates as of Dec '17:							
Assessed balance	-0.0504	0.0209	-0.0138	0.0058	-0.0118	0.0139	-0.0239
($M)	(0.0035)	(0.0028)	(0.0031)	(0.0021)	(0.0016)	(0.0016)	(0.0054)
Max module age	-0.0086	-0.002	-0.00001	-0.0006	-0.0035	-0.0022	-0.0158
(yrs)	(0.0002)	(0.0001)	(0.0002)	(0.0001)	(0.0001)	(0.0001)	(0.0003)
Share AB >9 yrs	0.006	-0.0097	-0.0074	-0.0018	0.0042	0.0027	-0.0049
(%)	(0.0020)	(0.0012)	(0.0018)	(0.0012)	(0.0008)	(0.0006)	(0.0030)
Unfiled returns	-0.0184	-0.0081	-0.0063	-0.0028	-0.003	-0.0019	-0.037
(1/0)	(0.0010)	(0.0006)	(0.0009)	(0.0005)	(0.0005)	(0.0005)	(0.0016)
Modules	0.0051	0.0035	0.0015	0.0008	-0.002	-0.0013	0.007
(# of non-Form 941)	(0.0003)	(0.0002)	(0.0002)	(0.0001)	(0.0001)	(0.0001)	(0.0004)
Modules	-0.0009	0.0017	0.0026	0.0009	-0.0012	-0.0006	0.0023
(# of Form 941)	(0.0002)	(0.0002)	(0.0002)	(0.0001)	(0.0001)	(0.0001)	(0.0004)
Constant	0.1613	0.0387	0.042	0.0283	0.0481	0.0284	0.3182
	(0.0035)	(0.0023)	(0.0026)	(0.0018)	(0.0018)	(0.0013)	(0.0049)
Income Dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes
SOA Dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Status Dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	266,890	266,890	266,890	266,890	266,890	266,890	266,890
Adjusted R^2	0.075	0.022	0.020	0.010	0.021	0.012	0.124
Mean dep. var.	0.076	0.028	0.040	0.017	0.016	0.008	0.172

Notes: heteroskedasticity-robust standard errors in parentheses. Certification is instrumented using SSN; SSN was used to randomly select a 5% control group that was held out from initial certification.

The results for the outcome of newly taking decertification actions are shown in Table 2.2. Consistent with the graphs above, we find that certification, instrumented using SSN, leads to increased activity by taxpayers in IAs, OICs, and CNC-Hardship designations. We do not find a significant effect on bankruptcies, nor for full resolution, either by payment or by abatement. These results suggest that certification does have a positive effect on compliance regarding tax debt, but perhaps not in the way that a casual observer might expect (i.e., by leading taxpayers to pay their balances in full). However, recall that the population of certified taxpayers includes a significant proportion of taxpayers with substantial debt that has not been resolved by other IRS actions over many years.

One concern with the passport program is that a taxpayer, once certified, could attempt to "game the system" by taking action that leads to decertification but does not materially affect their tax debt. For example, a taxpayer could take action to get decertified, use the time when initially decertified to make any necessary passport-related requests, and subsequently renege on any commitments that led to decertification. To test whether certification leads to more long-lasting compliance, we pair the outcome of newly taking each decertification action with the restriction that a taxpayer remain uncertified at the end of our observation window, in December 2019 – that is, we test whether certification leads taxpayers not just to take action to get decertified but also to remain uncertified. The results of these tests suggest that there is indeed a persistent effect on compliance; certification leads to a higher probability of new and maintained IAs and CNC designations. Appendix Table 2.9 shows the RCT IV regression results for tests of the combination outcome of taking a given compliance action and remaining uncertified at the end of our observation window. These specifications test whether the certification effect is long-lasting, and the results show that it is. Certification causes a significant positive increase in the probability of starting a new IA or being designated CNC and remaining uncertified at the end of the observation period. As further support for the certification effects being largely persistent, we observe that nearly all certified taxpayers who request an installment payment plan follow through with the process and remain in an approved IA. This is not a surprise, as the program is designed to discourage this sort of gaming behavior; a taxpayer who reneges on a prior agreement would be subsequently re-identified as eligible for certification, fall back into certification, and thus could have their passport revoked.

Finally, we consider several balance-related outcomes. Given the notched structure of certification, we do not expect a strong response of small payments. Once a taxpayer is certified, they must pay off or otherwise resolve all of their certified debt; merely paying off an amount to get below the certification threshold of $54,000 (lower in earlier years) does not lead to decertification, unless they are in an approved installment plan. Nevertheless, we test the effect of certification on payments (both in dollars and as a share of December 2017 assessed balance), and on changes in balance, changes in log balance, and binary variables indicating a fall in balance or increase in balance. The results for the balance-related outcomes are shown in Appendix Table 2.8.

We do find a small positive effect on the probability of payment, although this likely reflects, at least in part, the new IAs we saw before. Consistent with our expectations, we do not find a strong effect of certification on payments on average when measured in dollars, although we do find a marginally significant effect on payments as a share of starting balance, suggesting that certification causes 0.34 percentage points more to be paid off, on average, over the two years leading up to December 2019. We similarly find no significant effect on changes in balance, although interestingly we do find that certification makes it less likely that a taxpayer will have a fall in balance, and more likely that they have an increase in balance. One explanation for this is that the actions that certification induces – IAs, OICs, and CNCs – require a taxpayer to be current on all their current and prior tax filings. Certified taxpayers induced to take these compliance actions may thus have filed late returns with a balance due so they could qualify for an IA, OIC or CNC resolution. Turk et al. (2016) and Collins et al. (2018) found similar responses to the filing of tax liens. To test for this, we add several combination outcomes: taking an action (IA, OIC, or CNC) *and* adding new modules. Similarly, we add the combinations of *not* taking action and adding new modules, or *not* taking action and increasing assessed balance. The results provide support for the hypothesis that the increase in assessed balance is a side effect of certification driving IA, OIC, and CNC action: the positive effect of certification on adding new modules is only found among those also taking compliance actions. Appendix Table 2.10 reports the full results of this test.

The results of the RCT analysis suggest that there is indeed a positive effect of passport certification on tax debt compliance behavior. This is striking, given what we noted earlier about the eligible population including a mix of taxpayers, some of whom may not place much value on holding a passport, or may not anticipate needing to request one from the State Department in the near future. Furthermore, the results we find in the RCT analysis could be considered a lower bound on the effect of certification, as it is possible some of the control group taxpayers were aware that they were eligible for certification and acted in response to this. Although most of those in the control group were not certified, some may still have taken action to avoid the risk of future certification, which would provide a downward bias to the difference between the treatment and control groups as a causal estimate of the treatment.

2.5.2. Heterogeneity

Having considered the effect of certification overall, we now turn to several analyses which develop a more nuanced understanding of the effects of certification. We begin by noting that there may be a difference in program response for taxpayers with pre-existing debt relative to taxpayers with new debt. Recall that there was a large group of taxpayers with significant balances that would have been eligible for certification for years prior to 2018, if the program had been in place during those years. These taxpayers comprise the “pre-existing” delinquent taxpayers, distinct from the “new” delinquent taxpayers that become newly eligible once the program begins.[63] Understanding whether the response in these two groups differs is important because it is conceivable that the new taxpayers offer the best evidence for what effect the passport program will have going forward.

We begin with a graphical comparison, taking the same approach as above for the overall effect of certification, but now splitting the samples into preexisting and new taxpayers. There are clear differences in the base level of activity for the various actions we study, as shown in Appendix Figure 2.11. In general, the new cases are more likely to take action; this is to be expected, as these cases are being certified during the normal progression of the collection process. Thus, the certification is coming immediately after the triggering events for eligibility. The pre-existing cases are being certified well after the triggering events and in many instances after other collection actions have taken and the debt was not resolved.

Comparing the rate of activity between treatment and control SSNs, within the pre-existing and new groups, suggests that there may be differences in the treatment effect of certification for the two groups. We test for this using the same regression framework as above, now splitting the sample into the pre-existing and new cases. The results are shown in Appendix Table 2.11. Certification has a significant effect on the probability of new IAs and of any resolution, for both groups. The significant effect for new OICs and new CNCs is found only in the pre-existing cases. This is consistent with the notion that, for the pre-existing cases, it is more likely that the only

[63] We define as “pre-existing” observations those that had an eligible balance above $50,000 in December 2016 and that did not accrue additional modules between December 2016 and March 2018, when certifications began. These individuals would have been certified had the program been in operation in December 2016 but, as it was not, they were not certified until the rollout began. We define the remaining cases as “new” cases – these are cases that either had new modules accruing between December 2016 and March 2018, or did not have a sufficiently high eligible balance in December 2016 to warrant certification, but had such a balance during the passport program rollout.

viable resolution available is an administrative resolution like and OIC or CNC. The new cases are more likely to have the ability to pay and not be eligible for OIC or CNC resolutions.

We next examine whether the effect of certification is stronger among taxpayers who are more likely to have or want passports. Although we cannot determine who is a passport holder directly, at least for those taxpayers with recent tax filings we can observe taxpayer characteristics that are correlated with passport holding (based on an analysis of state-level passport holding (Florida 2011)). About one-third of the taxpayers in our main specification do not have recent tax filings and so are excluded from these tests. However, non-filing is not randomly distributed. Among the new cases, 88% have filings, while only 54% of pre-existing cases have filings. This means that, compared to the main results above, these heterogeneity tests are weighted more towards the effect on new cases. We investigate four characteristics that are correlated with passport holding: (1) income: higher-income individuals are more likely to hold passports; (2) age: older individuals are more likely to hold passports; (3) location: border state residents are more likely to hold passports; and (4) tax filing markers that indicate foreign activity.[64] As above, we start by comparing behavior graphically, splitting by treatment or control SSN and each of the passport proxy characteristics. We then incorporate these proxy characteristics into our regression analysis. Appendix Figure 2.12 compares the share of treatment and control groups newly taking any decertification action each month, splitting by the four different proxies: total positive income, age, border/non-border state, and foreign tax markers. We see clear differences in the activity levels when splitting by income (higher-income taxpayers are more likely to take action) and by foreign tax markers (those with

[64] Specifically, we identify foreign activity as any of the following: claiming the Foreign Tax Credit or the Foreign Earned Income Exclusion; filing from a foreign address; or filing a Foreign Bank Account Report (FBAR), for any tax year prior to and including 2017.

foreign activity are more likely to take action), but similar activity levels when splitting by age and border/non-border state.

Adding these treatment intensity proxies to the IV framework requires adding the characteristics as well as interactions of the characteristics with the certification dummy and with the SSN dummy. We use the characteristic X SSN interaction dummy as an instrument for the characteristic X certification dummy. Table 2.3 presents the results when testing the effect of certification and including these interactions on the binary outcomes of newly taking various actions (the same outcomes tested above in Table 2.2). We see a significant differential effect on the combined action outcome for the income interaction, but not for age, residing in a border state, or foreign activity. This appears to be driven by a stronger effect on IAs and full resolution by payment for those with higher incomes.[65] The marginally significant interaction effects for foreign tax activity on new OICs and new abatement could be explained by the relative sophistication and use of tax preparers by those with foreign activity, relative to those without.

[65] While reflecting a greater probability of passport holding, the stronger response among higher-income taxpayers may also reflect a greater ability-to-pay, especially when considering the new IA and full resolution outcomes. In the State Denied analysis below, in which all individuals had demonstrated a demand for passports, we find differences in response by income supportive of the ability-to-pay explanation: new IAs and full resolutions were more common for above-median income taxpayers, while new OICs and CNC-Hardship designations were more common for below-median income taxpayers.

Table 2.3: RCT IV treatment intensity regression results

	Taking new action any time Mar '18-Dec '19				Fully resolved as of Dec '19		Any of six listed actions
	IA (1)	OIC (2)	CNC (3)	Bankruptcy (4)	By payment (5)	By abatement (6)	(7)
Certified	0.0131	0.0082	0.0111	-0.0038	-0.0031	0.0001	0.0231
	(0.0063)	(0.0045)	(0.0055)	(0.0042)	(0.0027)	(0.0024)	(0.0096)
TPI > median TPI	0.0429	-0.0002	-0.024	-0.0061	0.0092	-0.0027	0.0218
	(0.0083)	(0.0055)	(0.0061)	(0.0045)	(0.0040)	(0.0026)	(0.0114)
Cert X TPI>median	**0.0236**	**-0.0052**	**-0.0083**	**0.0013**	**0.0088**	**0.0004**	**0.0202**
	(0.0085)	**(0.0056)**	**(0.0063)**	**(0.0046)**	**(0.0041)**	**(0.0027)**	**(0.0117)**
Age > median age	0.005	0.0026	0.0165	-0.0041	0.0009	-0.00001	0.0201
	(0.0080)	(0.0054)	(0.0061)	(0.0044)	(0.0038)	(0.0026)	(0.0112)
Cert X Age>median	**-0.0032**	**-0.0016**	**-0.0021**	**0.0024**	**0.0057**	**0.0012**	**0.0012**
	(0.0082)	**(0.0055)**	**(0.0063)**	**(0.0045)**	**(0.0039)**	**(0.0027)**	**(0.0115)**
Border state	0.0064	0.0025	0.0013	-0.0036	0.0085	0.0021	0.0166
	(0.0081)	(0.0054)	(0.0061)	(0.0044)	(0.0040)	(0.0027)	(0.0113)
Cert X Border state	**-0.0012**	**0.004**	**-0.0024**	**-0.0005**	**-0.0039**	**-0.0003**	**-0.0048**
	(0.0083)	**(0.0056)**	**(0.0063)**	**(0.0045)**	**(0.0041)**	**(0.0028)**	**(0.0116)**
Foreign tax filings	0.0168	0.0221	-0.0045	-0.0124	0.0142	-0.0045	0.0358
	(0.0171)	(0.0118)	(0.0109)	(0.0072)	(0.0103)	(0.0055)	(0.0226)
Cert X Foreign tax	**-0.008**	**-0.0223**	**0.0011**	**0.0097**	**0.0083**	**0.0104**	**-0.0047**
	(0.0176)	**(0.0120)**	**(0.0112)**	**(0.0074)**	**(0.0106)**	**(0.0057)**	**(0.0231)**
Dec '17 Covariates	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Income Dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes
SOA Dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Status Dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	179,813	179,813	179,813	179,813	179,813	179,813	179,813
Adjusted R^2	0.064	0.014	0.020	0.006	0.025	0.016	0.084
Mean dep. var.	0.1090	0.0410	0.0530	0.0240	0.0220	0.0100	0.2410

Notes: heteroskedasticity-robust standard errors in parentheses. Certification and certification-characteristic dummies are instrumented using SSN and SSN-characteristic dummies; SSN was used to randomly select a 5% control group that was held out from initial certification.

To further explore the interaction of income with the effect of certification, we identify the quartiles of total positive income among all 2017 tax filings, and categorize our taxpayers into each of these.[66] We then run the same interaction regressions as above, now using the TPI quartiles and their certification and SSN interactions. The results are shown in Appendix Table 2.12, and confirm that the certification effect is larger for those with higher income, concentrated in the top quartile of income. Appendix Figure 2.13 shows these results graphically. In general, we find that the effect of certification on actions that require some payment are higher for those with higher incomes (new IAs and full resolution by payment). We also find that the effect on the probability of any resolution increases with income, likely driven by new IAs being a large component of the overall resolution effect.

2.6. Effects of denied passport-related requests

As of January 2020, about 10,000 certified taxpayers had tried to do something passport-related and been denied (including new applications, renewals, and modifications).[67] This population offers another opportunity to study a direct effect of the passport program, as these are taxpayers who are certainly treated by the program (in contrast to certifications more generally, where taxpayers are only effectively treated when certified if they anticipate wanting to make passport-related requests in the near future). We refer to this set of taxpayers as the "State Denied" group.

As noted earlier, this is certainly a special group. These are individuals who were notified by the IRS that any passport request would be denied, and yet still chose to make a request. This group likely includes a mix of individuals. Some may not have received or read the certification notices,

[66] Our estimated quartile boundaries using a 1% random sample of filings are roughly $18,000, $40,000, and $82,000.

[67] The total number of certified taxpayers with a request denied by the State Department as of January 2020 was about 12,000, of which about 2,000 appear as secondary taxpayers with joint tax liabilities (spouses on joint tax returns). To avoid double counting, we exclude denials related to secondary taxpayers, and analyze only denials related to the taxpayer listed as the primary taxpayer.

so their responses are purely to the denied request. Previous literature studying mail notification experiments has found a wide range of letter reading rates, from as low as 20% to as high as 70% (Perez-Truglia and Cruces 2017, Bottan and Perez-Truglia, Betting on the House: Subjective Expectations and Market Choices 2020, Nathan, Perez-Truglia and Zentner, My Taxes are Too Darn High: Why do Houeholds Protest Their Taxes? 2021), although these rates are likely context-dependent and may not apply directly to our setting of enforcement letters sent by the IRS. Others in the group may have received and read the notices, but guessed that the IRS was bluffing, so their response is to both the denied request and to learning about the IRS' ability to enforce this new passport program. The mixed composition of this group needs to be kept in mind in interpreting the average response of this group to having a passport application denied.

We begin our analysis by showing the share of these taxpayers taking various actions in the months before and after their passport-related requests were denied by the State Department due to their certified status. One potential concern is that any graph of the share taking action, conditional on being certified in the prior month, may show some increase in such activity simply because, in order to be certified, one must *not* have taken any actions leading to decertification, and thus the probability of taking these actions in the following month must increase, or at least stay the same; we call this the mechanical effect (a similar pattern was seen above in Figure 2.2). To address this concern, we compare the behavior of the State Denied taxpayers to a control group. We take a random sample of taxpayer-month observations among currently certified taxpayers, matching the size of the State Denied group, and define the sampled month as t=0 for each of these taxpayers. We then observe these sampled individuals' behavior in the months before and after, and similarly calculate the share of this control group taking actions in each relative month. This group can still exhibit the mechanical effect (action in month t=1 may jump up simply because we condition on being certified in month t=0). Thus, comparing the behavior between the State Denied group (showing treatment + mechanical effect) and the control group (showing only the mechanical effect) allows us to see the effect of the denial of a passport-related request as the difference between the two groups.

Figure 2.3 compares the share of the State Denied group and the control group that were newly decertified in each month. We show six months before and twelve months after t=0 (the month in which a passport request was denied for the State Denied group, in orange, or the randomly selected month for the control group, in gray).[68] The figure shows that the denial of a passport-related request leads to an immediate jump in new decertifications for the State Denied group, above and beyond the mechanical effect that can be seen in the smaller jump for the control group. Between month t=0 and month t=6, the cumulative difference in new decertifications is about 20 percentage points (50% for the State Denied group vs. 30% for the control group).

Figure 2.3: Share getting newly decertified, State Denied vs. control group

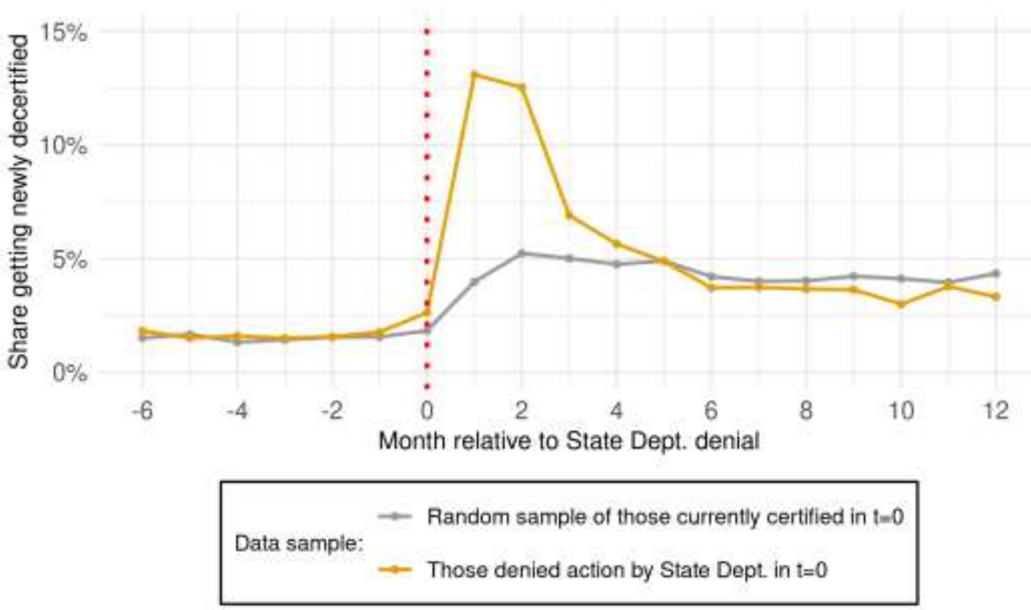

[68] We stop our outcome observations for this analysis in February 2020, to avoid any change in behavior due to COVID-19, which among other things undoubtedly reduced the immediate value of holding a passport due to travel restrictions. This means that not all taxpayers have twelve months of post-t=0 data. The count of taxpayers with available data, and thus who are included in the monthly share calculations, are shown in Appendix Figure 2.14.

Notes: This figure plots the share of each of two groups newly getting decertified in each month. In orange is the State Denied group, whose passport-related requests were denied in month t=0. In gray is the control group, a random sample of taxpayers who were certified as of month t=0.

We next consider separately the primary actions a taxpayer can take to get decertified. We first look to see whether taxpayers fully resolve their balance, by either payment, abatement, or other means.[69] Appendix Figure 2.15 compares the share of taxpayers in the State Denied and control groups resolving in full in the months relative to denial. The State Denied group is more likely to fully resolve both by payment and by abatement, with payment about five times as common as abatement. The two groups are roughly equally likely to fully resolve by other means. We also consider the other compliance actions we study above in the RCT analysis: IAs, OICs, CNCs, and Bankruptcies. Appendix Figure 2.16 compares the share newly taking each such action between the State Denied and control groups. The State Denied group is more likely to enter into new IAs, OICs, and CNC designations, with IAs about 10 times as likely as OICs, and twice as likely as CNCs. We can also split these samples further based on their recent income tax filings to understand more about the types of taxpayers taking each action. We find that those with above-median income are more likely to respond by starting a new IA, while those with below-median income are more likely to newly make OICs or enter CNC-Hardship status (see Appendix Figure 2.17).

These figures show the month-by-month differences in the share of each group taking action, with an initial large spike in activity by the State Denied group immediately after denial, followed by several months in which the State Denied group remains more active than the control group, with that difference tapering out over time. To estimate the cumulative effect of State Department

[69] Specifically, if a taxpayer's balance goes to zero and their payments total at least 90% of the starting balance, we classify this as a full resolution by payment. If abatements total at least 90% of the starting balance, we classify this as a full resolution by abatement. Taxpayers with balance going to zero that do not have large payments or abatements are classified as full resolution by other means.

denial, we can thus aggregate the difference in activity over the months after State Department denial. Table 2.4 compares the cumulative share of those newly taking each action over the months t=0 to t=6, demonstrating that State Department denial leads to significant increases in compliance behavior over the six months following denial.

Table 2.4: Cumulative share comparison, State Denied vs. control group

	Sum of share newly taking each action over months t=0 to t=6							
	Decertification	WTZ (Pay)	WTZ (Abate)	WTZ (Other)	IA	OIC	CNC-Hardship	Bankruptcy
State Denied group	49.4%	2.5%	0.5%	3.0%	13.6%	2.4%	6.4%	0.9%
Random sample	29.9%	0.5%	0.1%	2.5%	2.2%	1.2%	1.4%	0.6%
Difference	19.5%	2.0%	0.4%	0.4%	11.3%	1.2%	5.0%	0.3%

Notes: This table presents the cumulative share of taxpayers in each group taking each action over months t=0 to t=6. The State Denied group includes certified individuals who were denied a passport-related request in month t=0. The control group is a randomly selected group of taxpayers who were certified as of month t=0. The selected actions leading to decertification are not necessarily mutually exclusive, nor exhaustive of all actions that can lead to decertification. WTZ=Went To Zero balance.

Denial-induced compliance actions translate into making more payments, as well. Appendix Figure 2.18 shows the share of each group making positive net payments in each month. The share among the State Denied group is slightly above the control group in the pre-period, at around 15%. After denial, the State Denied share increases sharply, to about 20%, while the control group remains flat at 15%.[70] Comparing the balances of the two groups, taxpayers in the State Denied and control groups have similar starting average assessed balances, in the month prior to denial (t = -1), of about $200,000. The State Denied group then makes larger cumulative payments, on average, over the following months. The difference in cumulative payments between the State Denied and control groups shows the effect of denied passport requests. Over the six months post-

[70] If one looks only at changes in assessed balances, rather than at payments, the effect appears smaller because some State Denied taxpayers add balance-due modules *after* their denials, increasing their total assessed balance. As before in the RCT analysis, this is likely driven by the requirement that taxpayers be current on all tax filings before administrative resolutions such as IAs, OICs or CNCs are granted; to take compliance actions, taxpayers with unfiled returns for prior tax years must file those returns, which often result in new balance-due modules. Although this results in a higher assessed balance, it is a positive compliance outcome; becoming current on filing obligations is an important step towards becoming fully compliant in terms of their remittance obligations.

denial, the State Denied group paid an average $10,000, while the control group paid an average of $4,000, implying a denial effect of $6,000. Over the 12 months post-denial this difference is slightly higher, at $9,000.[71] Appendix Table 2.13 shows these comparisons in more detail. We also do these calculations separately for pre-existing and new cases, and find a larger effect of denial for the new cases: an incremental payment difference over 12 months of $12,600 for new cases vs. $7,300 for pre-existing cases.

We interpret these results as clear evidence that the threat, indeed the reality, of a passport denial or revocation induces positive compliance behavior for a non-trivial fraction of debtors. Two aspects should be noted. First, while about half of the State Denied group took an action to remove the passport restriction during the six months following the denied request, the other half did not take any such action. This suggests that the value of having a passport or renewed passport was worth less to the latter half than the most attractive avenue to decertification. Second, and related, all of these taxpayers were informed this could happen, and they applied for a passport or passport renewal anyway. As we noted earlier, this suggests they may not be entirely representative of the average taxpayer or tax debtor.

2.7. Policy implications

What does our empirical analysis imply about the revenue and welfare implications of the passport program? To answer this question, we consider the estimates we have produced of the effect of passport certification and the behavioral response of those denied passport-related requests. The results from our heterogeneity analyses above suggest that there are important differences between the pre-existing stock of eligible taxpayers and the taxpayers that became

[71] As above for the binary actions, we find a larger payment effect for those with above-median income: over the six-month period post-denial, the difference in cumulative payments for above-median income taxpayers is $7,800 vs. $1,500 for below-median income taxpayers.

newly eligible once the program was rolled out. We thus provide separate estimates for these two groups.

In Section 2.5, we investigated the behavioral effect of certification. We concluded that, as expected, the per-taxpayer effects were smaller than those in the State Denied analysis, but in many cases still statistically significant and non-trivial. We estimated a marginally significant effect on payments as a share of starting balance of 0.34% of the starting balance (a 4.3% increase in the payoff rate from a mean of 7.8%). Recall that our main specification excluded taxpayers with a balance above $1 million; when we include all taxpayers, the estimate is slightly lower at 0.33%. Applied to the average starting balance of about $195,000, this implies that certification causes an additional $644 in payments, over the year following certification. Again, this reflects a mix of new and pre-existing cases. Estimated separately, we find larger effects among new cases: 0.57% for new versus 0.11% for pre-existing. Applied to their average starting balances, this implies an additional $1,112 and $222 in payments, for new and pre-existing cases, respectively.

The State Denied analysis relies on a comparison of those certified taxpayers who were denied some request by the State Department to a control group of randomly-selected certified taxpayers. By comparing their behavior in the months following the passport request denial, we learn about the incremental effect of denial, above and beyond the base level of activity for certified taxpayers over the same time. In Section 2.6, we concluded that for the average passport-denied taxpayer, denial led to additional payments of about $9,000 over the twelve months following denial. This $9,000 estimate reflects a mix of new and pre-existing cases. When we do the same analysis separately for these groups, we find that the effect on new cases is stronger than for pre-existing cases. Denied requests lead to average additional payments of $12,600 for new cases, and $7,300 for pre-existing cases.

Table 2.5 summarizes these estimated effects. The separate estimates for pre-existing and new cases are useful in evaluating the passport program. The pre-existing cases represent a one-time boost in revenue; these cases would have been certified earlier had the program been in place all along, but in fact were all certified in a short period of time when the program was rolled out. The new cases better represent the ongoing effect of the program; these cases become newly eligible over time, and so offer a better measure of how the program may affect collections going forward. Note also that these estimates likely undercount the total effect, because some of the response to certification or denial includes taxpayers initiating Installment Agreements, whose effect on payments would be realized over a longer period. Finally, keep in mind that some of the impacts of the program may not result in new payments but are still productive outcomes for tax administration, such as new CNC designations and new abatements.

Table 2.5: Marginal revenue estimates

		Pre-existing cases	New cases
Effect of certification			
Coefficient estimate	*% of balance*	0.11%	0.57%
Average total balance	*$*	$202,000	$195,000
Estimated payment effect	*$/person*	$222	$1,112
Effect of denied requests			
Estimated payment effect	*$/person*	$7,300	$12,600

Notes: This table reports the marginal revenue estimates, separately for pre-existing and new cases. Certification effects come from the RCT IV approach in Section 2.5; denied request effects come from the cumulative payment comparison in Section 2.6.

As has been noted by Slemrod and Yitzhaki (1987) and Keen and Slemrod (2017), that the program raises net revenue is not dispositive as to whether it had a positive welfare impact. A welfare analysis must consider that the additional money raised is not a resource gain but is rather a transfer from private citizens (albeit those with tax debt) to the government, which has social value (only) to the extent that the social value of the government services it enables exceeds the social value of the foregone taxpayer income; in addition, a welfare analysis should consider not just the additional dollars collected, but also the administrative cost of the program and the private costs induced by it. For a monetary sanction, the private costs are predominantly compliance costs, but in the case of a collateral sanction they also include the utility loss of those whose passport is denied. The marginal administrative costs should be net of any resource cost saving due to running a somewhat less extensive passport program due to the tax-related certifications and revocations.

The program passes the welfare test if the following condition holds (where ϕ is the marginal social value of public spending, assumed to be greater than one):

$$\phi(\Delta Revenue - \Delta Administrative\ Cost) - \Delta Private\ Cost > 0$$

As noted above, we estimate that for new cases, certification leads to an additional $1,112 in payments and denial to an additional $12,600 in payments. We observe that 3-4% of certified taxpayers subsequently had requests denied by the State Department, so that the total effect can be approximated as $1,112 + $12,600*0.035 = $1,553.[72] For the IRS-related administrative costs, after an initial fixed cost of setting up the program, the marginal cost of an additional certification includes the time of IRS and State Department staff in processing the certification, and the cost of mailing letters from the IRS and State Department; based on internal IRS estimates, we assume these together amount to $25 per certified taxpayer. Then, so long as the private cost of certification is less than $\phi(\$1{,}553 - \$25)$, we can conclude that the passport program is welfare-improving

[72] The effects of denial and certification are distinct and so can be added without resulting in double counting. Our denial estimates rely on a control group of certified taxpayers, so that the estimated effects are above and beyond any certification effect. In addition, those taxpayers who are denied requests by the State Department must not have responded to certification itself, or else they would have been decertified and no longer subject to denial. This interpretation is confirmed by a robustness check of our certification regression analysis that excludes the State Denied taxpayers, and finds consistent results (see Appendix Figure 2.9).

within this framework.[73] As noted above, the private cost of collateral sanctions should include the foregone utility from restricted travel options due to passport denial.

This calculation does not incorporate any saving in administrative costs due to the fact that more labor-intensive enforcement policies do not have to be applied, so that the welfare gain of the certification program may be understated. It also ignores the fact that, in principle, collateral sanctions could lower the cost of providing the service to which they restrict access. In this setting, for example, by restricting access to passports the program may make it less costly for the State Department to produce and monitor passports because fewer individuals request them. This effect could offset some of the administrative cost of administering the collateral sanctions, and if large enough could in principle result in net negative administrative costs. It also ignores the private utility costs of foregoing a passport and the benefits that having a passport provides.

2.8. Conclusions

This paper provides the first evidence on the effects of a large-scale collateral sanction program, studying the introduction and first two years of operation of the IRS' passport certification and revocation program. By leveraging an RCT-style implementation during program rollout and observing the behavior of a subset of seriously delinquent taxpayers who were denied passport requests, we provide evidence that the passport program leads a substantial number of taxpayers to take compliance actions they otherwise would not have. Our work suggests that tax agencies (and legislatures) can consider collateral sanctions as a viable option for improving tax compliance.

[73] The corresponding test using the estimates for the pre-existing cases would be that $\phi(\$465 - \$25)$ exceeds the private cost of certification for the pre-existing case taxpayers.

2.9. Appendix A: Additional figures and tables

Table 2.6 provides summary statistics about the population of first-time certified taxpayers, as of the time of our study. This includes all first-time certified taxpayers from when the passport program was rolled out in 2018 through April 2020, when certifications were paused due to COVID-19. The distribution of tax debt is skewed, so that although the average certified taxpayer had a total certified balance of about $197,000, the median was about half that, or $98,000. These balances are typically the result of several years of unpaid tax liabilities, with the median certified taxpayer having four modules (tax years) certified. The median annual income of those certified was about $60,000 (Adjusted Gross Income), again reflecting a skewed distribution in which the average is nearly twice as large, at $103,000.

Table 2.6: Summary statistics for first-time certified taxpayers

	Mean	St Dev	25th Pctile	Median	75th Pctile	# Obs
Certified balance						
Assessed balance, penalties, and interest ($ thousands)	$197	$1,146	$68	$98	$172	394,000
Number of modules	5	4	2	4	7	394,000
Age of oldest module (years)	7	3	4	7	9	394,000
Most recent tax filing prior to certification						
Total positive income ($ thousands)	$149	$5,551	$30	$68	$134	293,000
Adjusted gross income ($ thousands)	$103	$1,926	$24	$60	$120	293,000
Primary income source						
Wages	0.18	0.39	0.00	0.00	0.00	394,000
Schedule C income	0.28	0.45	0.00	0.00	1.00	394,000
Interest, dividends, and capital gains	0.25	0.44	0.00	0.00	1.00	394,000
Age in 2017 (years)	53	11	46	53	61	379,000

Notes: Values are rounded for disclosure purposes. Includes first-time certifications through April 2020, when certifications were paused due to COVID-19.

Figure 2.4 shows the monthly count of certifications from program rollout in 2018 through April 2020. The initial months show the rollout of the program, with July 2018 the single largest month. This reflects the fact that there was a stock of eligible taxpayers who would have been certified in prior years, if the program had been running during those years. We thus observe a large initial set of certifications during the months of program rollout, followed by a steady flow of new certifications. The figure also shows that some taxpayers cycle in and out of certification; in orange are certifications of taxpayers who were previously decertified. We focus our analysis in this paper on initial certifications.

Figure 2.4: Monthly count of certifications (first-time and subsequent instances)

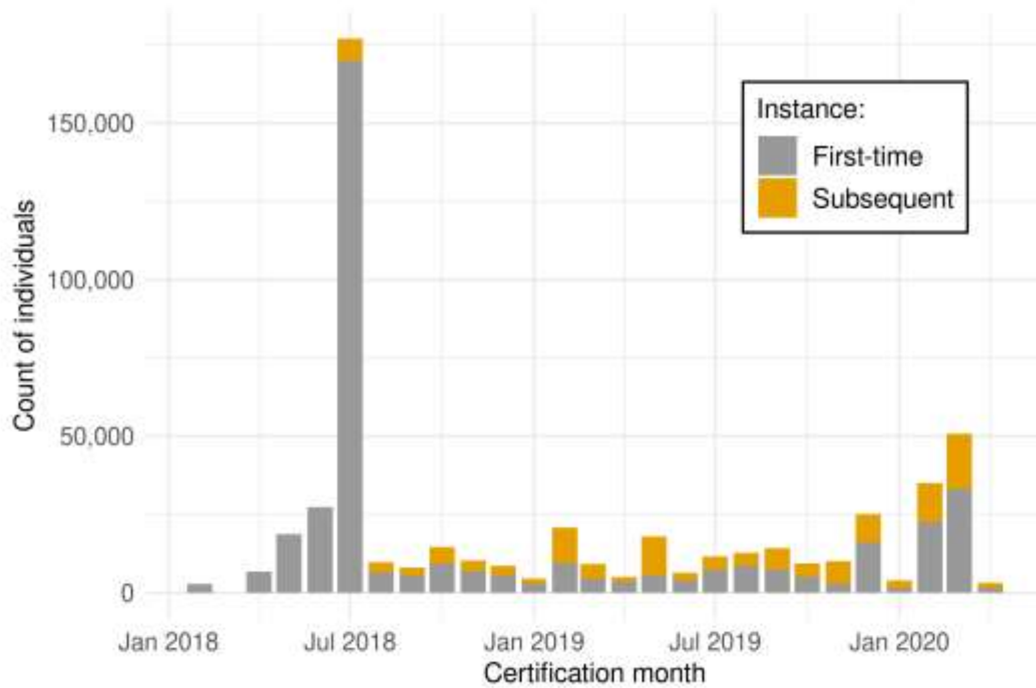

Notes: This figure plots the monthly count of individuals having tax debt newly certified under the passport program. In gray are individuals with tax debt certified for the first time; in orange are individuals with tax debt certified for a second or subsequent time.

Figure 2.5 shows the monthly count of decertifications. There is some indication of seasonality in decertifications, with December 2018 and 2019 showing much higher numbers, relative to January 2019 and 2020, respectively.

Figure 2.5: Monthly count of decertifications (first-time and subsequent instances)

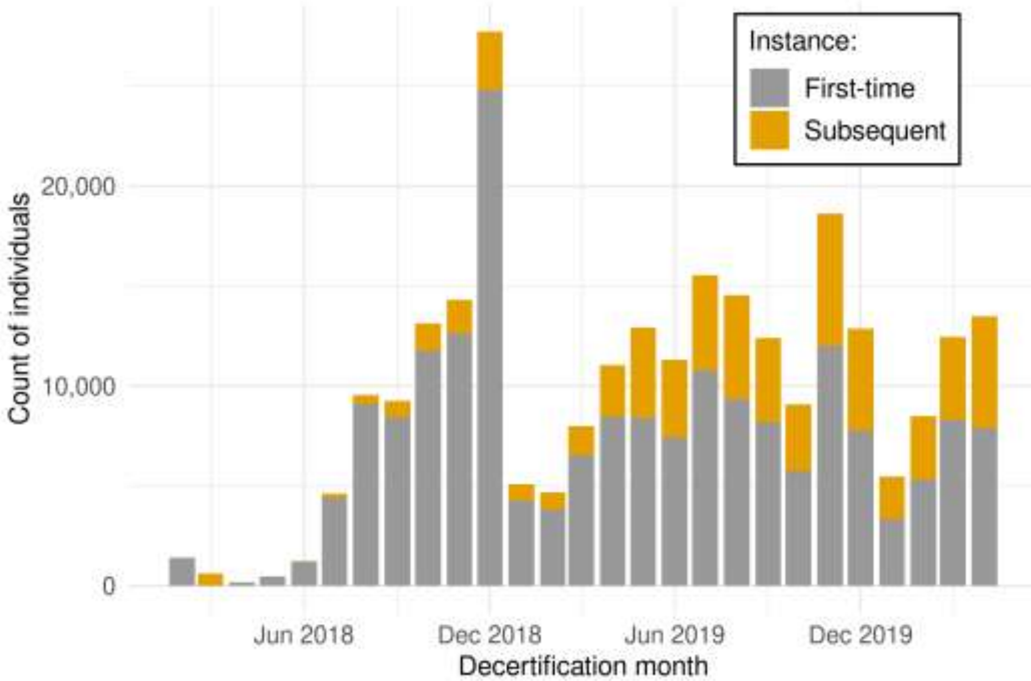

Notes: This figure plots the monthly count of individuals having their tax debt certification under the passport program reversed. In gray are individuals with tax debt decertified for the first time; in orange are individuals with tax debt decertified for a second- or subsequent time.

Figure 2.6 shows the daily count of visits to the IRS' webpage providing information about the passport program. The figure shows that there was an initial spike in visits to the page when it was published, coincident with a number of news articles describing the new program. There is also a spike in visits in July 2018, when the majority of initial certifications occurred. This provides some evidence that at least some certified taxpayers make an effort to learn more about the program once they are certified.

Figure 2.6: Daily page views for IRS passport program webpage

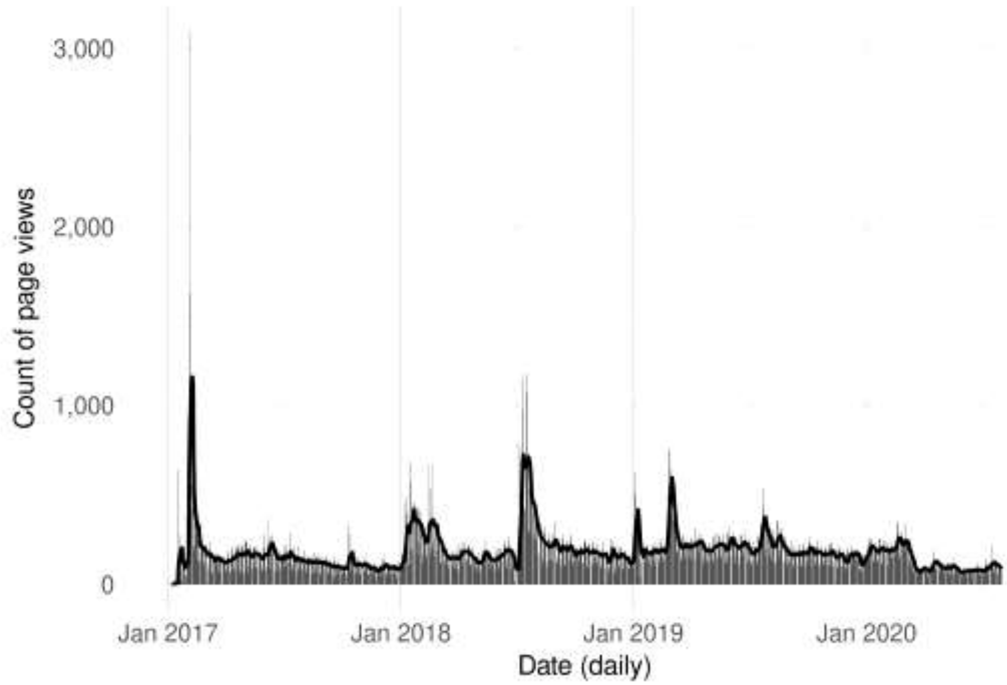

Notes: This figure plots the daily page visits to the IRS webpage describing the passport certification and revocation program. The gray bars present the daily visits; the black line presents the prior-seven-day rolling average.

Figure 2.7 compares the treatment and control SSN groups from the RCT analysis on each of four individual compliance actions. The top panel compares these groups by the monthly share newly taking each action: new Installment Agreements (IAs), new Offers-in-Comprise (OICs), new designation of Currently Not Collectible (CNC) due to Hardship, and new Bankruptcies. The bottom panel compares the groups by the share currently in each status, by month. The figure shows that certification led to noticeably more IAs and CNCs, as well as a smaller increase in OICs.

Figure 2.7: RCT analysis, comparison only by SSN groups, four actions leading to decertification

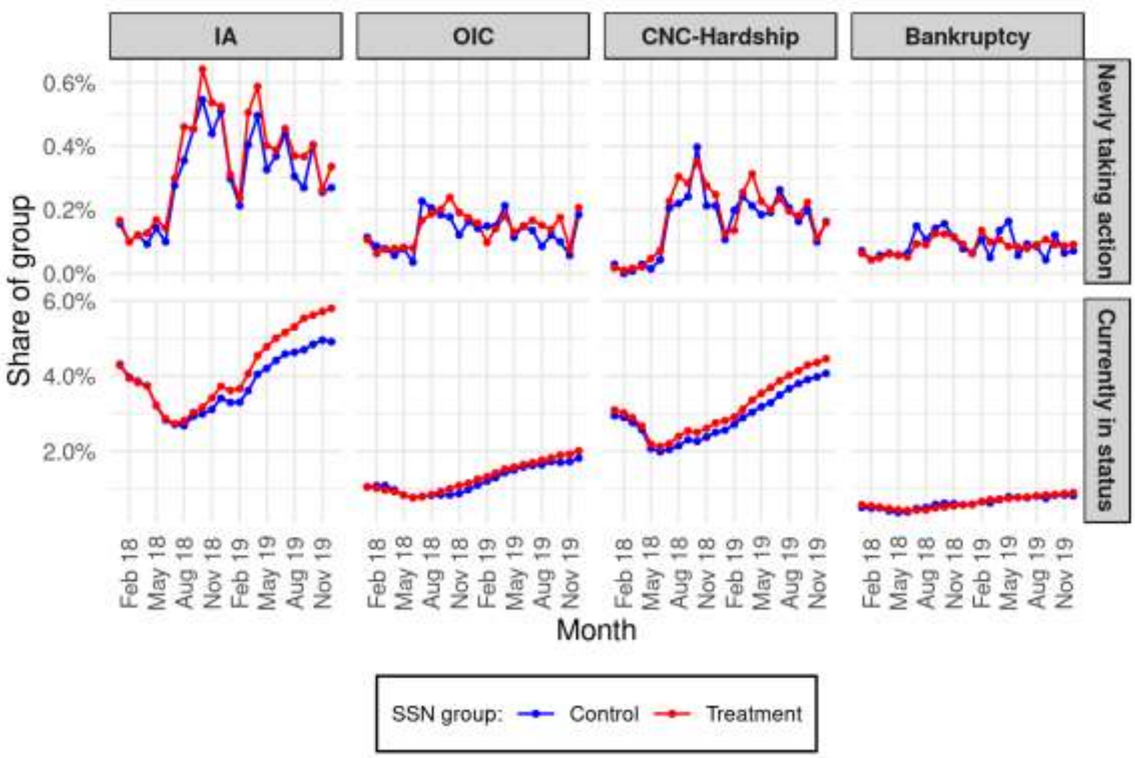

Notes: This figure compares the share newly taking each action (top panel) or currently in a given status (bottom panel) for taxpayers who by SSN were assigned to treatment or control groups during the passport program rollout. All of the treatment group was certified, while approximately 20% of the control group migrated into treatment during the twelve-month RCT phase. The difference between the two groups thus reveals the effect of certification.

Figure 2.8 compares the treatment and control SSN groups from the RCT analysis on the cumulative share who have fully resolved their balances (top left panel) and by the share newly going to zero balance by three means: full payment, abatement, or other means. The figure shows that, unlike for the actions above, there is no clear difference between the groups on their likelihood of full resolution. This reflects the higher cost of full resolution as a compliance action, relative to the less costly options of new IAs, OICs, CNCs, and Bankruptcies shown above.

Figure 2.8: RCT analysis, comparison only by SSN groups, full resolution

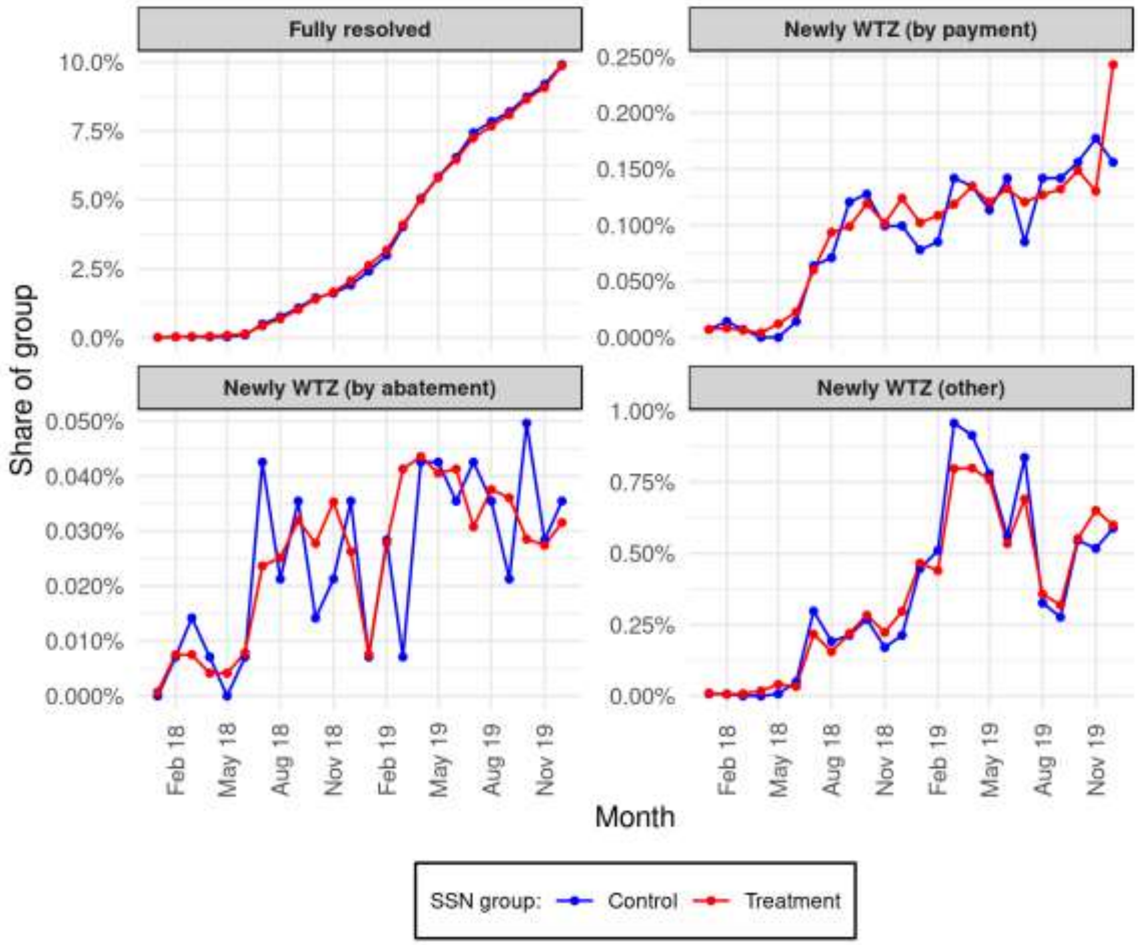

Notes: This figure compares treatment and control taxpayers by the share making a positive net payment, the cumulative share that has fully resolved over time, and the share newly having balance go to zero by payment or by other means.

Table 2.7 presents summary statistics for the data underlying the RCT analysis. Because of the filters applied (in particular, restricting to those with balance < \$1M), the average assessed balance among this population is about \$145,000, which is smaller than the average of \$197,000 for the full population of certified taxpayers shown above in Table 2.6. The medians are still close, however, reflecting the skewed distribution of debts.

Table 2.7: RCT analysis, summary statistics for data used in main analysis

Statistic	Mean	St. Dev.	Min	Pctl(25)	Median	Pctl(75)	Max
Treatment SSN	0.95	0.22	0	1	1	1	1
Certified	0.96	0.20	0	1	1	1	1
Dec '17 Covariates							
Assessed balance ($K)	145.107	138.206	0.005	66.815	95.217	162.005	999.599
Max module age (yrs)	6.377	2.783	0.000	4.216	6.899	8.720	11.997
Share of assessed balance older than 9 years	0.075	0.210	0.000	0.000	0.000	0.000	1.000
Indicator for presence of unfiled returns	0.205	0.404	0.000	0.000	0.000	0.000	1.000
Number of annual modules	3.944	2.440	0.000	2.000	4.000	6.000	10.000
Number of quarterly modules	0.528	2.420	0.000	0.000	0.000	0.000	40.000
Balance outcomes (as of Dec '19)							
Assessed balance ($K)	130.29	171.19	0.0	52.7	84.5	151.1	15,013.7
Change in assessed balance ($K)	-14.82	129.16	-996.5	-41.0	0.0	15.4	14,747.5
Change in log assessed balance	-0.618	1.593	-6.905	-0.489	0.000	0.137	8.737
Fall in assessed balance	0.457	0.498	0	0	0	1	1
Increase in assessed balance	0.420	0.494	0	0	0	1	1
Cumulative payments (Dec '17 to Dec '19)							
Made payment	0.477	0.499	0	0	0	1	1
Net payment ($K)	12.128	55.241	-957.0	0.0	0.0	4.6	7,151.5
Payment as share of Dec '17 AB (%)	0.078	0.206	0	0	0	0	1
Decertification actions, ever taken over period from Mar '18 to Dec '19							
New IA	0.076	0.265	0	0	0	0	1
New OIC	0.028	0.166	0	0	0	0	1
New CNC	0.040	0.195	0	0	0	0	1
New Bankruptcy	0.017	0.130	0	0	0	0	1
Assessed balance to zero (by payment)	0.016	0.125	0	0	0	0	1
Assessed balance to zero (by abatement)	0.008	0.090	0	0	0	0	1
Assessed balance to zero (other)	0.076	0.265	0	0	0	0	1
Any resolution action other than AB zero (other)	0.172	0.378	0	0	0	0	1

Notes: all variables have 266,890 observations. Percentiles are rounded for disclosure purposes.

Table 2.8 shows the RCT IV regression results for several balance-related outcomes. As discussed in the text, the results suggest certification has a small positive effect on the probability of making a payment, and positive but less significant effects on the amount of payment made. The balance results are mixed and, as discussed in the text, reflect that the positive effect on compliance actions such as new IAs and new CNCs also induces filing of new modules. This can raise the assessed balance, but is still a positive compliance outcome as it is a necessary step for a taxpayer to resolve their balance.

Table 2.8: RCT IV analysis, results for balance-related outcomes

	Cumulative payments			Assessed balance, Dec '19 vs. Dec '17			
	Made payment (1)	Amount ($K) (1)	Share (of Dec '17 AB) (2)	d(bal) ($K) (3)	d(logbal+1) (4)	1(Decrease in AB) (5)	1(Increase in AB) (6)
Certified	**0.0193**	**0.1274**	**0.0034**	**0.9946**	**0.0121**	**-0.0108**	**0.0214**
	(0.0050)	(0.5893)	(0.0021)	(1.2138)	(0.0154)	(0.0050)	(0.0048)
Covariates as of Dec '17:							
Assessed balance	0.0733	57.2689	-0.057	-165.1757	-0.7703	0.0643	-0.0479
($M)	(0.0069)	(2.4481)	(0.0026)	(4.8166)	(0.0263)	(0.0070)	(0.0066)
Max module age	-0.0199	-2.4287	-0.0128	-8.0908	-0.1464	0.0517	-0.0501
(yrs)	(0.0004)	(0.0566)	(0.0002)	(0.1231)	(0.0015)	(0.0004)	(0.0004)
Share AB >9 yrs	0.0151	3.1728	0.0156	-86.0288	-1.9004	0.4531	-0.2552
(%)	(0.0047)	(0.3154)	(0.0013)	(1.2793)	(0.0180)	(0.0037)	(0.0035)
Unfiled returns	-0.0386	-2.2489	-0.0066	4.3581	0.1465	-0.0271	0.0518
(1/0)	(0.0025)	(0.2118)	(0.0008)	(0.7554)	(0.0075)	(0.0023)	(0.0023)
Modules	0.0124	-0.8404	-0.0034	4.0132	0.1589	-0.0135	0.0256
(# of non-Form 941)	(0.0005)	(0.0685)	(0.0002)	(0.1413)	(0.0016)	(0.0005)	(0.0005)
Modules	0.0088	-1.0601	-0.0045	1.6792	0.05	-0.0022	0.0076
(# of Form 941)	(0.0004)	(0.0471)	(0.0002)	(0.1187)	(0.0012)	(0.0004)	(0.0004)
Constant	0.5454	25.0678	0.2012	41.8917	-0.2329	0.2166	0.5715
	(0.0061)	(0.7868)	(0.0028)	(1.7023)	(0.0193)	(0.0063)	(0.0060)
Income Dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes
SOA Dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Status Dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	266,890	266,890	266,890	266,890	266,890	266,890	266,890
Adjusted R^2	0.169	0.073	0.124	0.115	0.218	0.165	0.208
Mean dep. var.	0.4770	12.13	0.0780	-14.8190	-0.6180	0.4570	0.4200

Notes: heteroskedasticity-robust standard errors in parentheses.

Figure 2.9 shows the RCT IV coefficient estimates under various alternative data specifications. The figure demonstrates that, in general, the estimates are not sensitive to these specifications.

Figure 2.9: RCT IV coefficient estimates using alternative data filters

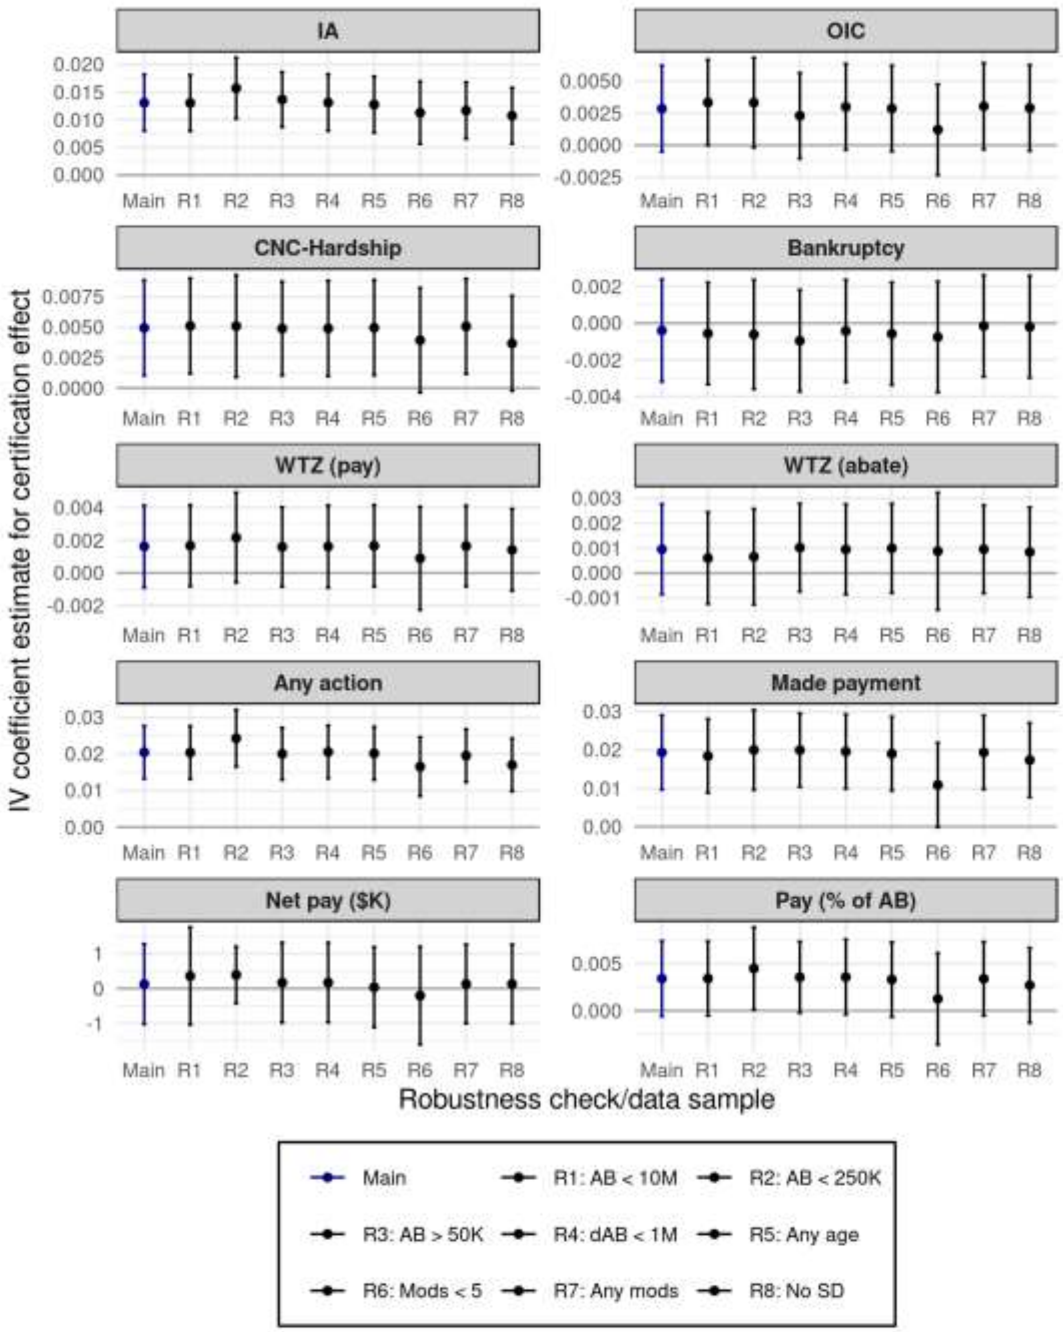

Notes: This figure presents the IV coefficient estimates, and 95% confidence intervals, when running the RCT IV analyses using alternative data filters. The main specification restricts to those with December 2017 assessed balance < $1M, with max module age < 12 years, and number of annual modules < 10 and quarterly modules < 40. R1-R3 adjust on December 2017 assessed balance, expanding up to $10M, restricting to those <$250K, or restricting to those >$50K, respectively. R4 restricts on the observed change in assessed balance, excluding any with a change >$1M. R5 removes the restriction on max module age. R6 and R7 adjust the module count restriction, with R6 restricting to those with <5 annual or <20 quarterly modules, and R7 removing the module count restriction. R8 removes taxpayers that have passport-related requests denied by the State Department, to confirm that the certification effects are not simply reflecting the effects of these denials.

Figure 2.10 shows the RCT IV coefficient estimates, using alternative end months. Our main specification uses December 2019 as the end month. The figure shows that the effect grows over time. We use December 2019 as our end month to allow sufficient time for taxpayers to take action after certification, but end the series prior to 2020 because the certification of "complex debt" cases began in early 2020. Many of the control group taxpayers were certified as part of the resulting update to the eligibility algorithm.

Figure 2.10: RCT IV coefficient estimates using different end months

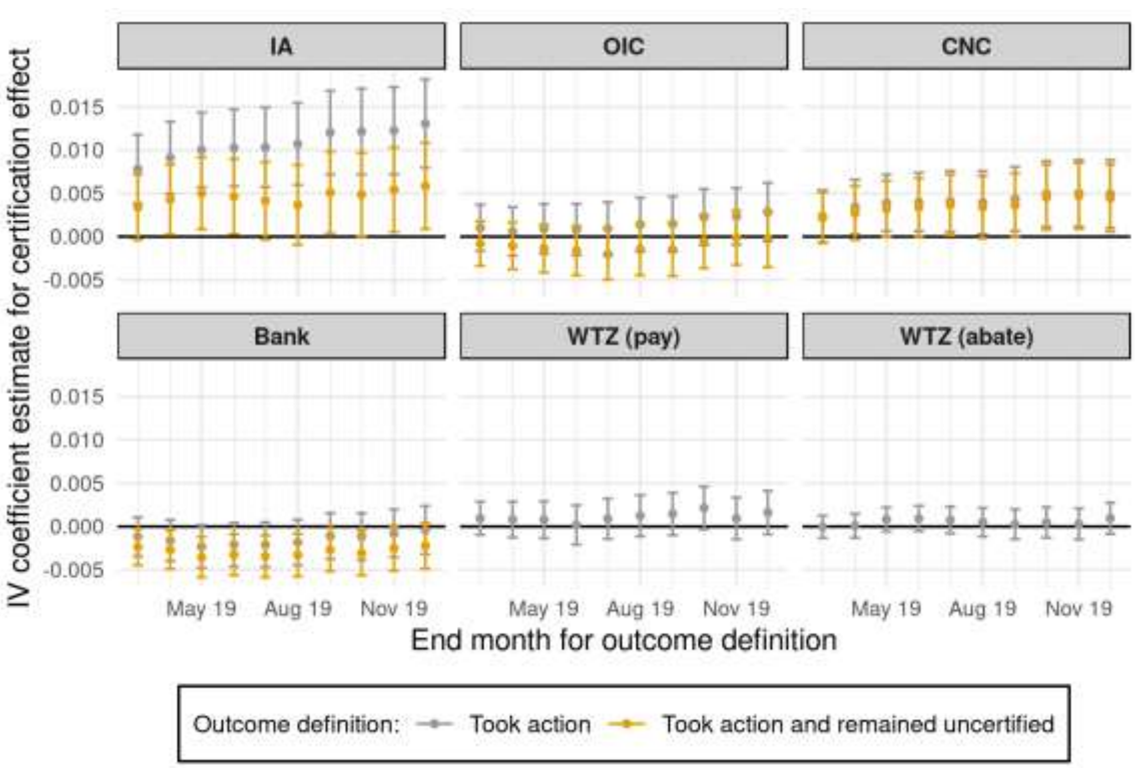

Notes: This figure presents the IV coefficient estimates, and 95% confidence intervals, when running the RCT IV analyses using different end months. Our main specification uses December 2019 as the end month.

Table 2.9 shows the RCT IV regression results for tests of the combination outcome of taking a given compliance action and remaining uncertified at the end of our observation window. These specifications test whether the certification effect is long-lasting, and the results suggest it is. Certification causes a significant positive increase in the probability of starting a new IA or being designated CNC, and remaining uncertified at the end of the observation period.

Table 2.9: RCT IV analysis, results for taking actions and remaining uncertified

	Taking new action and remaining uncertified as of Dec '19			
	IA (1)	OIC (2)	CNC (3)	Bankruptcy (4)
Certified	**0.0058**	**-0.0003**	**0.0045**	**-0.0022**
	(0.0025)	(0.0017)	(0.0020)	(0.0014)
Covariates as of Dec '17:				
Assessed balance	-0.0494	0.016	-0.0154	0.0018
($M)	(0.0033)	(0.0026)	(0.0030)	(0.0018)
Max module age	-0.0076	-0.0018	-0.00002	-0.0005
(yrs)	(0.0002)	(0.0001)	(0.0002)	(0.0001)
Share AB >9 yrs	0.009	-0.0078	-0.0069	-0.0011
(%)	(0.0020)	(0.0011)	(0.0018)	(0.0011)
Unfiled returns	-0.0172	-0.0066	-0.006	-0.003
(1/0)	(0.0009)	(0.0006)	(0.0009)	(0.0005)
Modules	0.0034	0.0027	0.0014	0.0006
(# of non-Form 941)	(0.0003)	(0.0002)	(0.0002)	(0.0001)
Modules	-0.0012	0.0014	0.0024	0.0002
(# of Form 941)	(0.0002)	(0.0002)	(0.0002)	(0.0001)
Constant	0.155	0.0374	0.0418	0.0261
	(0.0034)	(0.0022)	(0.0026)	(0.0017)
Income Dummies	Yes	Yes	Yes	Yes
SOA Dummies	Yes	Yes	Yes	Yes
Status Dummies	Yes	Yes	Yes	Yes
Observations	266,890	266,890	266,890	266,890
Adjusted R^2	0.065	0.017	0.019	0.009
Mean dep. var.	0.066	0.024	0.039	0.014

Notes: heteroskedasticity-robust standard errors in parentheses.

Table 2.10 shows the results from additional tests reconciling the findings that certification causes both an increase in compliance actions (IAs, OICs, and CNCs) and an increase in balance due. The table shows that those taking compliance actions in response to certification are also more likely to add new modules; a necessary condition for an IA, OIC, or CNC is to be current on all tax filings, so that some taxpayers may need to file previously unfiled returns in order to be eligible for taking these actions. The regression results show that certification leads to additional modules only when paired with taking a compliance action.

Table 2.10: *RCT IV analysis, testing explanation for increase in AB result*

	Combination outcomes: Take action *and* add modules			
New IA/OIC/CNC:	Yes	No	Yes	No
New modules for tax year:	<= 2016	<= 2016	>= 2017	>= 2017
	(1)	(2)	(3)	(4)
Certified	**0.0056**	**-0.0034**	**0.0104**	**-0.0047**
	(0.0020)	(0.0034)	(0.0027)	(0.0039)
Covariates as of Dec '17:				
Assessed balance ($M)	-0.0187	0.0447	-0.0268	0.0447
	(0.0030)	(0.0049)	(0.0039)	(0.0053)
Max module age (yrs)	-0.0026	-0.0129	-0.0086	-0.0139
	(0.0002)	(0.0003)	(0.0002)	(0.0003)
Share AB >9 yrs (%)	-0.0071	0.0106	-0.0051	0.0107
	(0.0015)	(0.0027)	(0.0020)	(0.0032)
Unfiled returns (1/0)	0.0173	0.0407	-0.0131	-0.0293
	(0.0011)	(0.0018)	(0.0010)	(0.0016)
Modules (# of non-Form 941	0.0008	0.0076	0.011	0.0225
	(0.0002)	(0.0004)	(0.0003)	(0.0004)
Modules (# of Form 941)	0.0018	0.0031	0.0001	0.0032
	(0.0002)	(0.0003)	(0.0002)	(0.0004)
Constant	0.0421	0.1117	0.1424	0.1872
	(0.0026)	(0.0043)	(0.0036)	(0.0050)
Income Dummies	Yes	Yes	Yes	Yes
SOA Dummies	Yes	Yes	Yes	Yes
Status Dummies	Yes	Yes	Yes	Yes
Observations	266,890	266,890	266,890	266,890
Adjusted R^2	0.020	0.020	0.083	0.123
Mean dep. var.	0.04	0.1100	0.08	0.1770

Notes: heteroskedasticity-robust standard errors in parentheses. The results suggest that the addition of new modules is driven by those taking new compliance actions (IA, OIC, and CNC), and that this helps explain why certification makes it more likely for a taxpayer to have an increase in assessed balance, on average.

Figure 2.11 compares taxpayers when splitting both on SSN groups (treatment and control) and by whether the case is pre-existing or new. We define as “pre-existing” observations those that had an eligible balance above $50,000 in December 2016 and that did not accrue additional modules between December 2016 and March 2018, when certifications began. These individuals would have been certified had the program been in operation in December 2016 but, as it was not, they were not certified until the rollout began. We define the remaining cases as “new” cases – these are cases that either had new modules accruing between December 2016 and March 2018, or did not have a sufficiently high eligible balance in December 2016 to warrant certification, but had such a balance during the passport program rollout. The figure shows that new cases have a higher level of each of the activities, and that there may be differences between the treatment effects for pre-existing and new cases. We investigate these quantitatively using the RCT IV regression framework.

Figure 2.11: RCT analysis, comparison only by SSN groups, split by Pre-existing and New cases

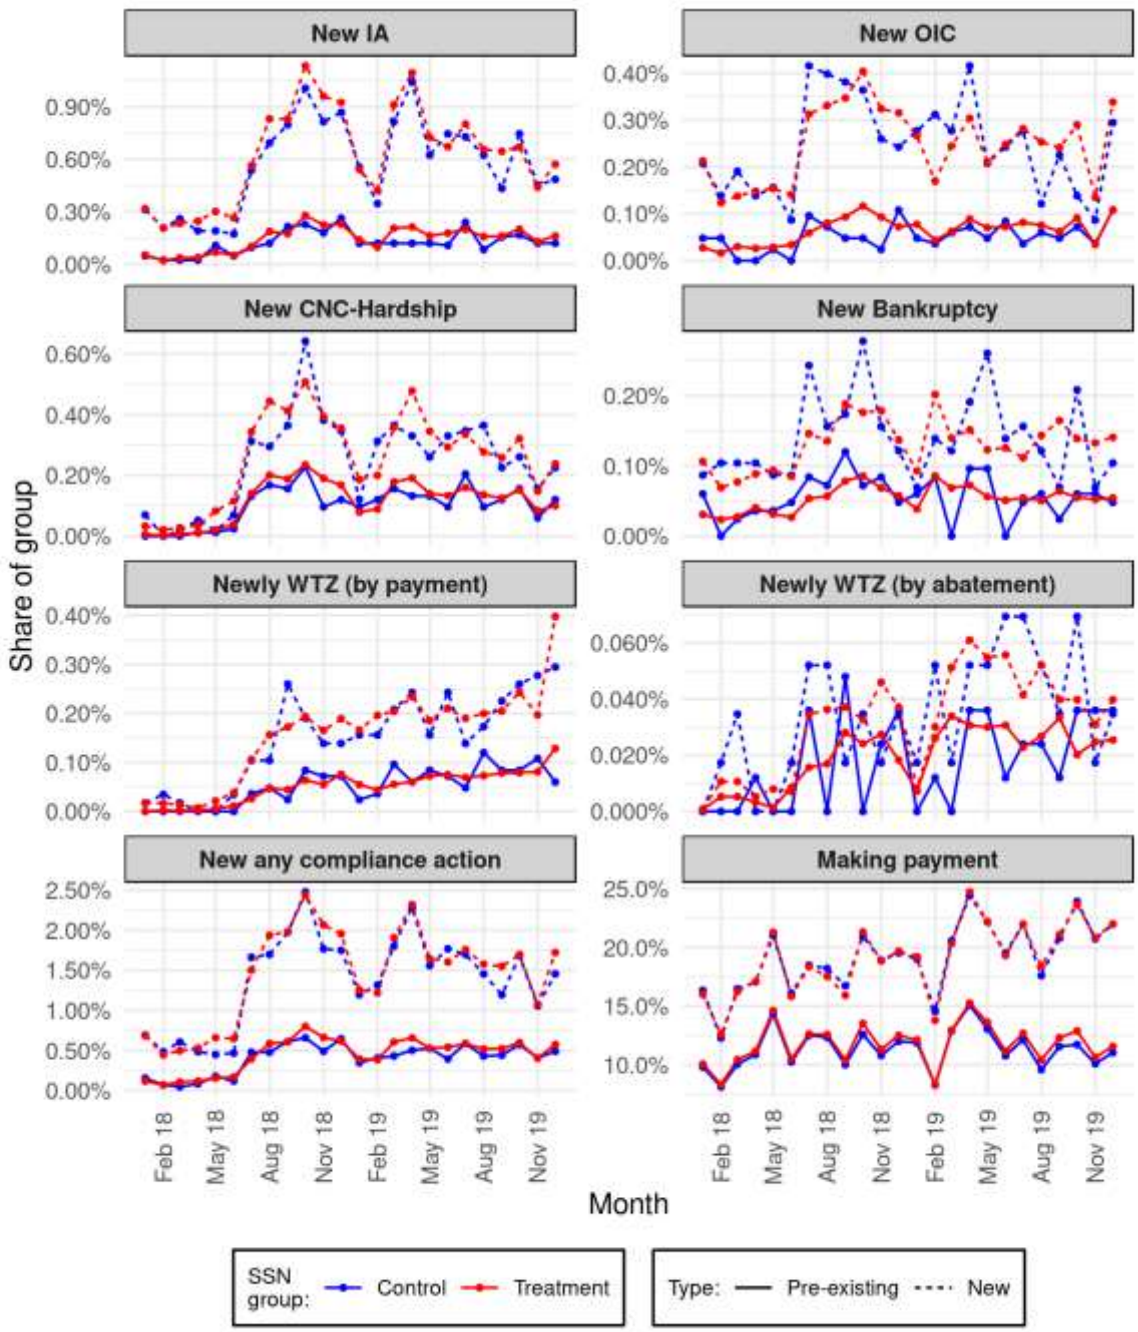

Notes: This figure compares treatment and control taxpayers by the share newly taking any of the binary actions leading to decertification, splitting the sample into taxpayers categorized as "pre-existing" or "new".

Table 2.11 presents the RCT IV regression results, testing for the effect of certification separately on pre-existing and new cases. Certification has a significant effect on the probability of new IAs and of any resolution, for both groups. Notably, the significant effect for new OICs and new CNCs is found only in the pre-existing cases, which is consistent with the notion that, for the pre-existing cases, it is more likely that the only viable resolution available is an administrative resolution like and OIC or CNC. The new cases are more likely to have the ability to pay and not be eligible for OIC or CNC resolutions.

Table 2.11: RCT IV analysis, results for separate analysis of Pre-existing and New cases

	Taking new action any time Mar '18-Dec '19				Fully resolved as of Dec '19		Any of six listed actions
	IA (1)	OIC (2)	CNC (3)	Bankruptcy (4)	By payment (5)	By abatement (6)	(7)
Panel A: Pre-existing debt							
Certified	**0.0055**	**0.003**	**0.0042**	**0.0006**	**-0.0003**	**0.001**	**0.0136**
	(0.0019)	(0.0012)	(0.0018)	(0.0012)	(0.0011)	(0.0009)	(0.0031)
Mean dep. var.	0.0320	0.0130	0.0260	0.0110	0.0080	0.0060	0.0910
Observations	155,317	155,317	155,317	155,317	155,317	155,317	155,317
Adjusted R^2	0.039	0.013	0.016	0.008	0.011	0.009	0.071
Panel B: New debt							
Certified	**0.0271**	**0.0018**	**0.0056**	**-0.003**	**0.0055**	**0.0009**	**0.0311**
	(0.0073)	(0.0049)	(0.0052)	(0.0038)	(0.0034)	(0.0023)	(0.0097)
Mean dep. var.	0.1370	0.0490	0.0580	0.0260	0.0270	0.0110	0.2860
Observations	111,573	111,573	111,573	111,573	111,573	111,573	111,573
Adjusted R^2	0.059	0.016	0.017	0.007	0.020	0.017	0.088
Dec '17 Covariates	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Income Dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes
SOA Dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Status Dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: heteroskedasticity-robust standard errors in parentheses.

Figure 2.12 compares taxpayers when splitting both on SSN groups (treatment and control) and by four separate characteristics that are correlated with passport-holding. Higher-income, older, border-state resident, and foreign-active taxpayers are apparently more likely to be holders of passports. The figure shows that there are level differences in activity when splitting by income and foreign activity. Table 2.3 in the main text reports the results of regression analysis studying these passport proxy characteristics.

Figure 2.12: RCT analysis, comparison only by SSN groups, heterogeneity in treatment effect

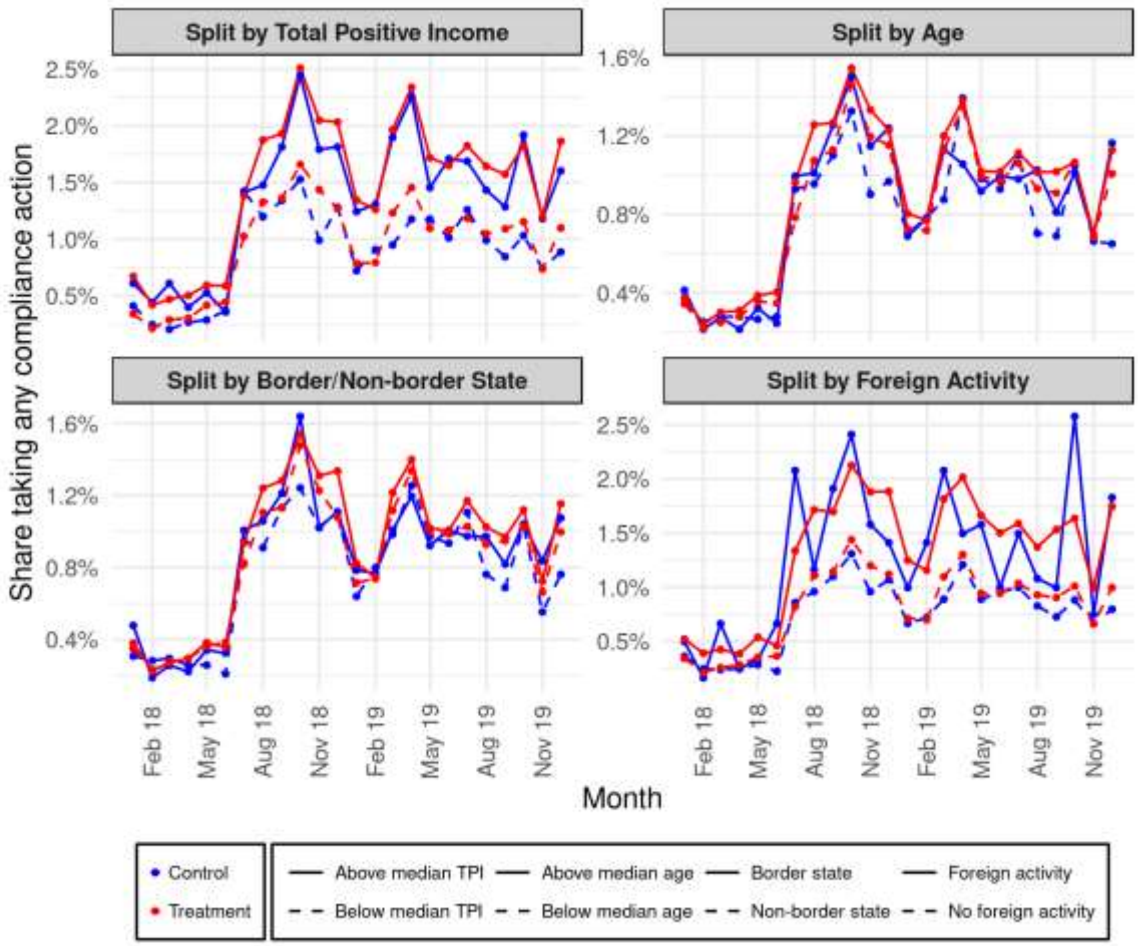

Notes: This figure compares treatment and control taxpayers by the share newly taking any of the binary actions leading to decertification, splitting on four proxies for passport holding: income, age, border/non-border state, and tax filing markers that indicate foreign activity.

Table 2.12 reports the RCT IV regression results, testing the interaction with income separately by quartiles. The results show that positive treatment effects are concentrated among those in the top quartile of income.

Table 2.12: RCT IV analysis, results with Total Positive Income interactions

	Taking new action any time Mar '18-Dec '19				Fully resolved as of Dec '19		Any of six listed actions
	IA (1)	OIC (2)	CNC (3)	Bankruptcy (4)	By payment (5)	By abatement (6)	(7)
Certified	0.0021	0.0103	0.011	-0.0045	-0.004	0.0024	0.0149
	(0.0058)	(0.0048)	(0.0068)	(0.0046)	(0.0032)	(0.0032)	(0.0105)
TPI in Q2	-0.0019	0.0113	-0.0028	-0.002	0.0004	0.0029	0.0132
	(0.0086)	(0.0075)	(0.0096)	(0.0065)	(0.0044)	(0.0045)	(0.0153)
Cert X TPI in Q2	**0.014**	**-0.0099**	**0.0007**	**0.0027**	**0.0002**	**-0.0033**	**0.0006**
	(0.0088)	**(0.0078)**	**(0.0099)**	**(0.0067)**	**(0.0045)**	**(0.0046)**	**(0.0157)**
TPI in Q3	0.02	0.0058	-0.0027	-0.0041	0.0003	0.001	0.0218
	(0.0094)	(0.0071)	(0.0092)	(0.0062)	(0.0044)	(0.0040)	(0.0149)
Cert X TPI in Q3	**0.0176**	**-0.0023**	**-0.0091**	**0.0038**	**0.0032**	**-0.0015**	**0.0121**
	(0.0096)	**(0.0073)**	**(0.0094)**	**(0.0064)**	**(0.0045)**	**(0.0041)**	**(0.0153)**
TPI in Q4	0.057	0.0009	-0.0323	-0.0122	0.0089	-0.0044	0.0248
	(0.0104)	(0.0069)	(0.0082)	(0.0059)	(0.0055)	(0.0038)	(0.0149)
Cert X TPI in Q4	**0.0344**	**-0.0083**	**-0.0087**	**0.0051**	**0.0163**	**0.0009**	**0.0363**
	(0.0107)	**(0.0070)**	**(0.0083)**	**(0.0060)**	**(0.0056)**	**(0.0039)**	**(0.0153)**
Dec '17 Covariates	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Income Dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes
SOA Dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Status Dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	180,989	180,989	180,989	180,989	180,989	180,989	180,989
Adjusted R^2	0.066	0.014	0.020	0.006	0.024	0.017	0.082
Mean dep. var.	0.1090	0.0410	0.0530	0.0240	0.0220	0.0100	0.2410

Notes: heteroskedasticity-robust standard errors in parentheses.

Figure 2.13 shows the RCT IV coefficient estimates on the income-quartile X certification interactions. The figure shows that, in general, the effect of certification increases with income. This is especially true for the IA and full resolution by payment outcomes.

Figure 2.13: RCT IV analysis, coefficient estimates for TPI quartile interactions

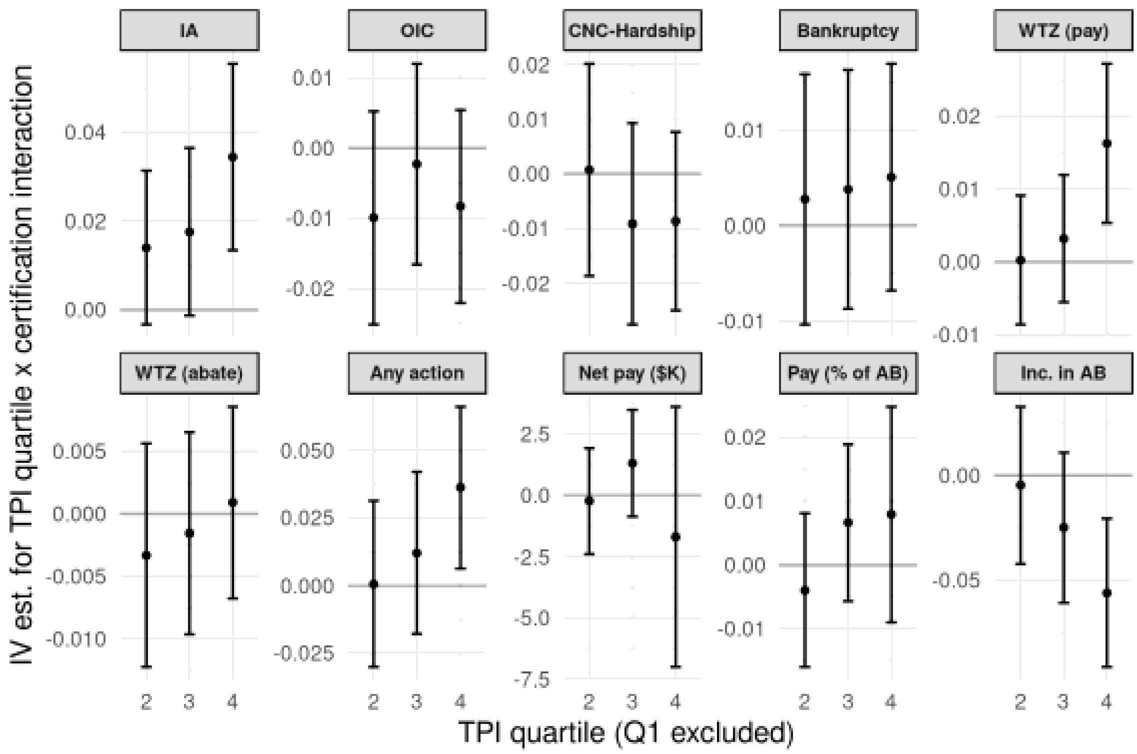

Notes: This figure plots the estimated treatment interaction coefficients and 95% confidence intervals, for TPI quartiles 2, 3, and 4 (quartile 1 is excluded), for ten separate outcomes.

Figure 2.14 shows the count of observations used for each relative month of the State Denied analysis. The observation counts decrease because we include only observations up through February 2020, to avoid any effect of COVID-19. So, for instance, a taxpayer whose request was denied in December 2019 would only have two months of post-denial observations.

Figure 2.14: Observation counts by relative month for State Denied analysis

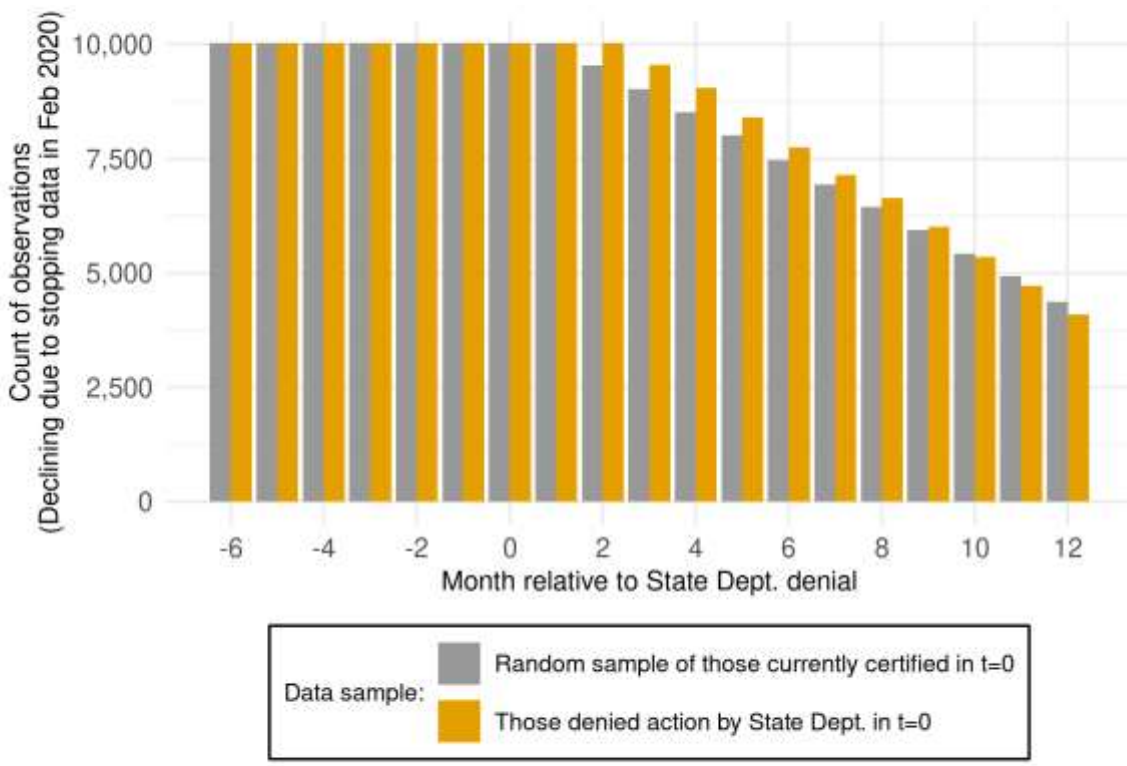

Notes: We only include outcome observations through February 2020, to avoid any effect of COVID-19. We include all denied passport requests through December 2019. This means that some taxpayers do not have a full twelve months of post-denial observations, and the size of the groups included for each relative month share calculation declines over time.

Figure 2.15 compares the State Denied and control groups on their monthly share newly resolving in full (i.e., balance going to zero) by payment, by abatement, or by any other means. Denial leads to an immediate increase in both full payments and abatements, although payments are about five times as common.

Figure 2.15: Share fully resolving balances, State Denied vs. control group

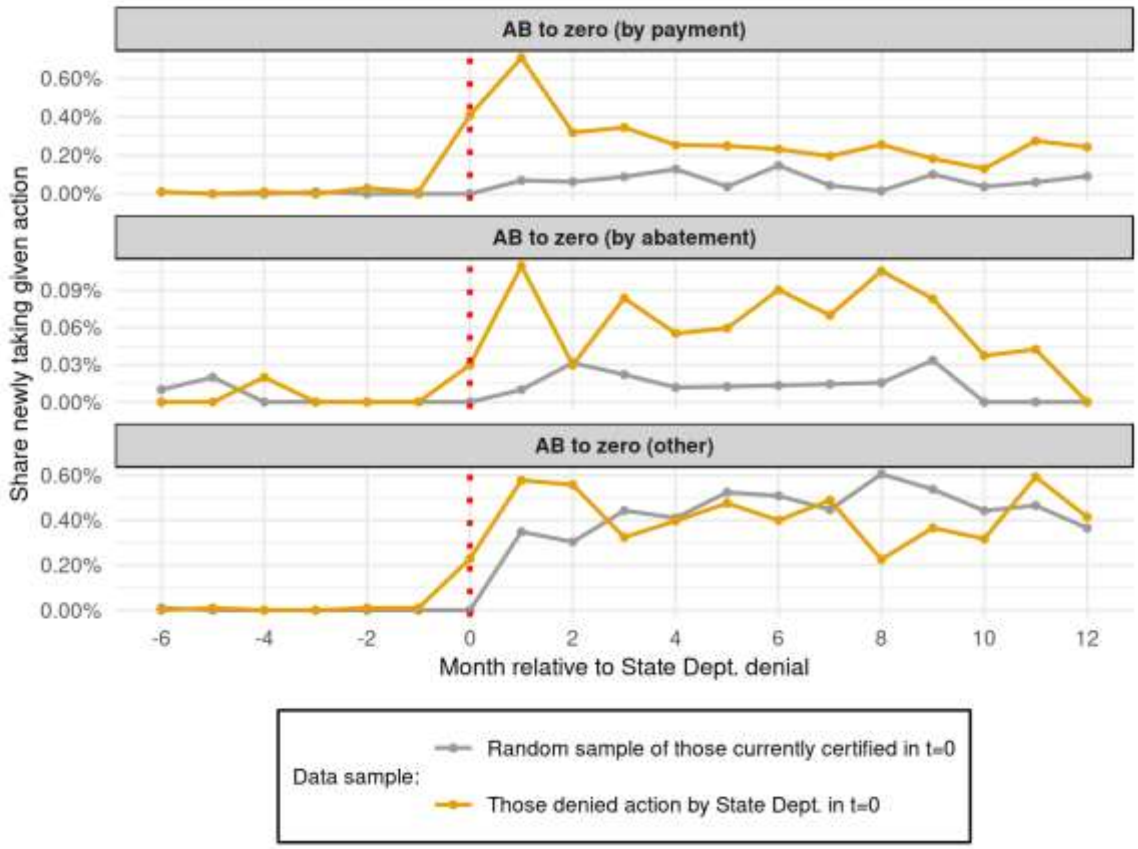

Notes: This figure plots the share of each of two groups newly resolving their balance in full, either by payment, abatement, or other means. In orange is the State Denied group, whose passport-related requests were denied in month t=0. In gray is the control group, a random sample of taxpayers who were certified as of month t=0.

Figure 2.16 compares the State Denied and control groups on their monthly share newly taking each of four compliance actions: IAs, OICs, CNCs, and Bankruptcies. Denial leads to more of each of these actions, though IAs and CNCs are much more common as responses, relative to OICs and Bankruptcies.

Figure 2.16: Share taking new actions leading to decertification, State Denied vs. control group

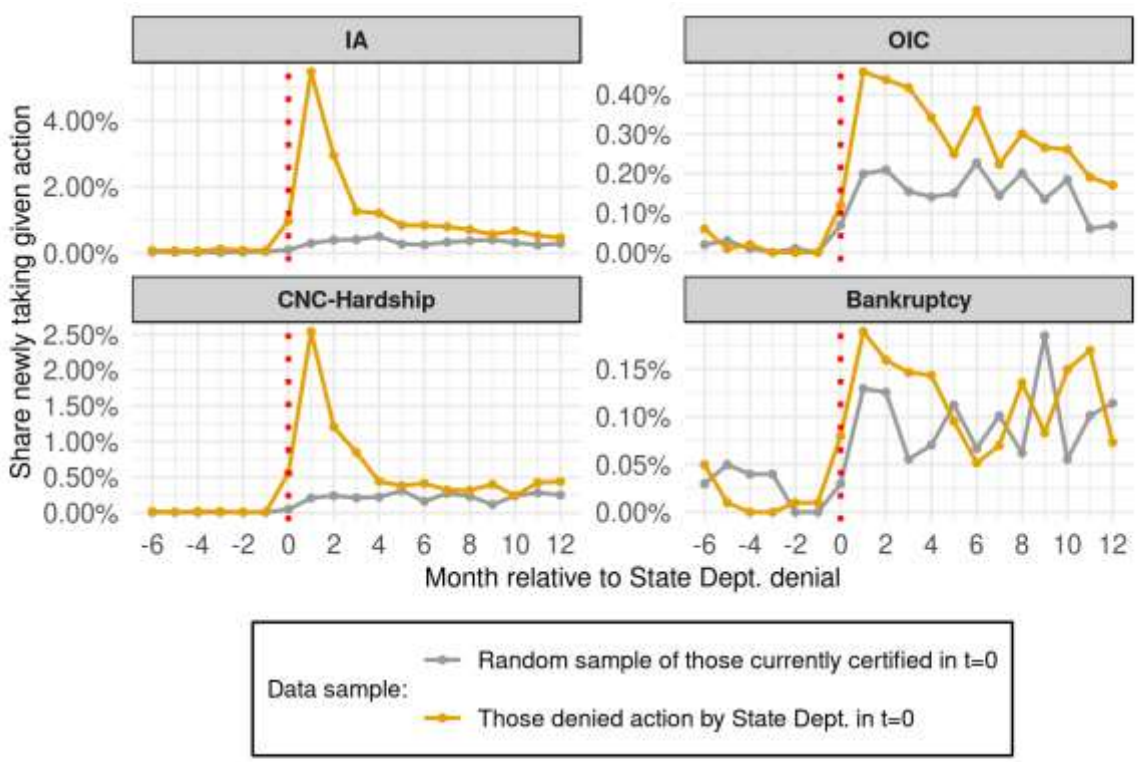

Notes: This figure plots the share of each of two groups newly entering into each status in each month, for the main actions that lead to decertification. In orange is the State Denied group, whose passport-related requests were denied in month t=0. In gray is the control group, a random sample of taxpayers who were certified as of month t=0. These actions are not necessarily mutually exclusive, nor exhaustive of all actions that can lead to decertification.

Figure 2.17 compares the State Denied and control groups on the share of each group currently in each of the four action statuses (IA, OIC, CNC, and Bankruptcy), and separately for those with above and below median income. The figure shows that the IA response is more common for higher-income taxpayers, while OICs, CNCs, and Bankruptcies are more common for lower-income taxpayers.

Figure 2.17: Share currently in certain statuses, State Denied vs. control group, split by income

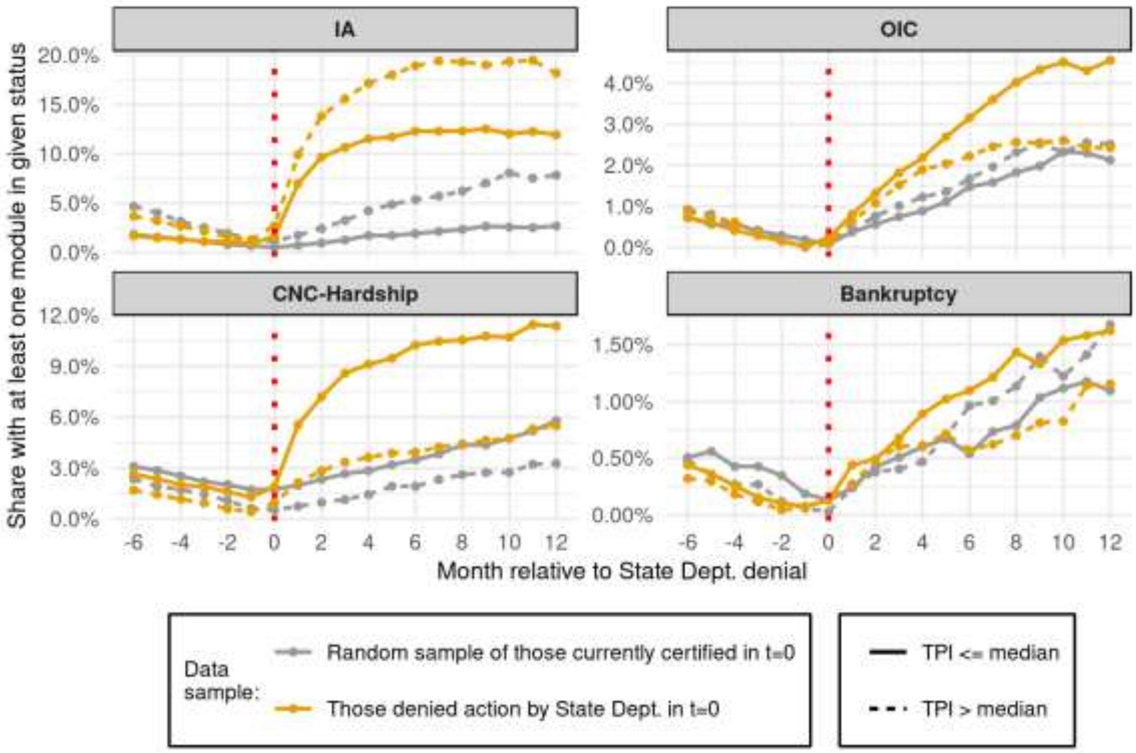

Notes: This figure presents the share of taxpayers in each group with at least one module in the given status. Taxpayers are split into the State Denied and control groups, and further into those with their most recent tax filing's total positive income above or below the median for the full set of taxpayers. Those without a recent tax filing are excluded.

Figure 2.18 compares the State Denied and control groups on their monthly share making any positive net payment (left panel) and adding new modules (right panel). There is a large increase in the share making payments, an increase of about five percentage points from a base of 15%. The addition of new modules (i.e., filing new returns with balance due) is consistent with the finding that denial leads to new actions like IAs, OICs, and CNCs, because these actions require taxpayers to be current on all tax filings, and thus taxpayers may need to file previously unfiled returns prior to taking these actions.

Figure 2.18: Share making payments and adding new modules, State Denied vs. control group

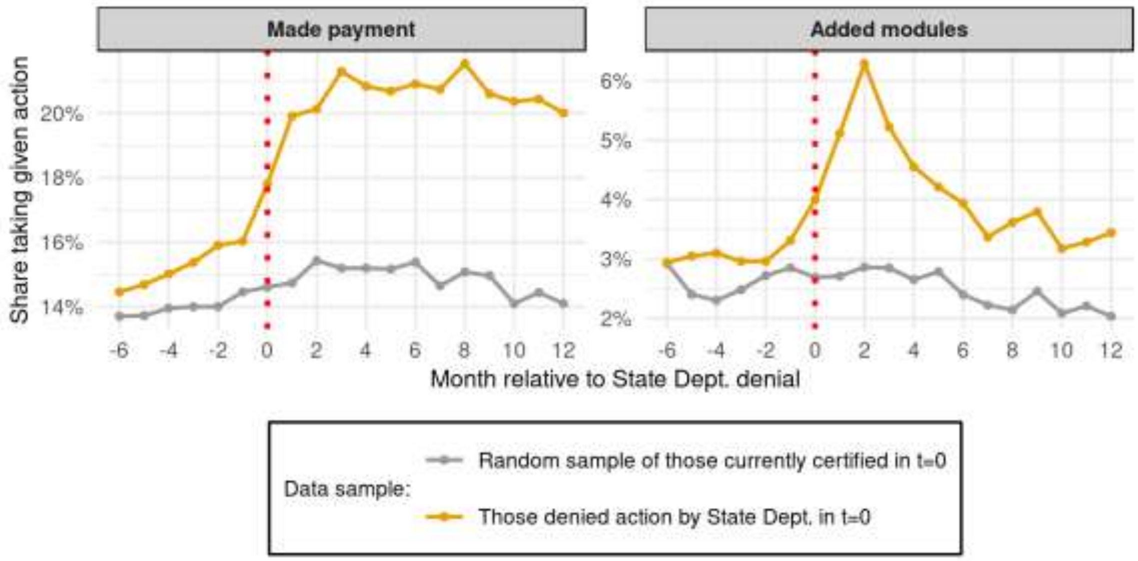

Notes: This figure presents the share of taxpayers in each group that make a payment each month. The high base share of payments each month is likely due to levies and other automatic payments.

Table 2.13 compares balances, payments, and abatements between the State Denied and control groups. The left two columns compare among all taxpayers (either State Denied or control group) whose final month is within our observation window for this analysis, through February 2020. The right two columns restrict to just those who do not add new modules, which reconciles the finding that assessed balances fall by less, on average, in the State Denied group; as noted above, some of the induced compliance actions require first filing previously unfiled returns, which may result in additional modules and increased balance due. The incremental payment effect of a denied request, over six months, can be found by taking the average payment of $10,000 among the State Denied group, and subtracting the average payment of $4,000 among the control group, resulting in an incremental effect of $6,000. For the twelve months post-denial, this effect is $9,000 ($15,000-$6,000). When we do the same 12-month calculation separately for pre-existing and new cases, we find that denial has a larger incremental payment effect on new cases: a difference of $12,600 for new cases vs. $7,300 for pre-existing cases.

Table 2.13: *Balance and cumulative payment comparison, State Denied vs. control group*

	Among all with end-month data		Subset to those not adding new modules	
	State Denied	Random sample	State Denied	Random sample
Panel A: 6-month cumulative outcomes				
Average assessed balance in t = -1 ($K)	$195	$206	$196	$208
Average assessed balance in t = 6 ($K)	$187	$194	$178	$190
% fall in total assessed balance	4.3%	6.0%	9.2%	8.5%
Average net payment from t = 0 to t = 6 ($K)	$10	$4	$11	$4
Average abatement from t = 0 to t = 6 ($K)	$6	$2	$6	$2
Panel B: 12-month cumulative outcomes				
Average assessed balance in t = -1 ($K)	$201	$190	$201	$186
Average assessed balance in t = 12 ($K)	$186	$171	$170	$158
% fall in total assessed balance	7.3%	9.7%	15.2%	15.1%
Average net payment from t = 0 to t = 12 ($K)	$15	$6	$17	$6
Average abatement from t = 0 to t = 12 ($K)	$9	$4	$9	$4

Notes: This table presents average values for taxpayers in the State Denied and control group taxpayers with data available in month t=6 (Panel A) or t=12 (Panel B), to ensure a direct comparison of changes in balance and payments. The left two columns include all such taxpayers; the right two columns restrict the comparison to only those without additional modules, to address the fact that some State Denied taxpayers were induced to add modules in order to take other compliance actions. All dollar figures are in thousands.

2.10. Appendix B: Indirect effects of the passport program

This paper focuses on the direct effects of the passport program. Specifically, we study what happens once an eligible taxpayer is certified (i.e., notified that a passport request will be denied absent appropriate action) and then, for the subset of certified taxpayers who make passport-related requests to the State Department and are denied, what happens after those denials. It is important to note, however, that the passport program could also have indirect effects on tax compliance, including on the behavior of those who already have some outstanding debt, as well as more generally on whether taxpayers incur debt at all.

In this Appendix, we address two potential indirect effects. First, we note that the program is designed with a notch: those with debt one dollar above the threshold are eligible for certification, while those just below are not. This suggests that one indirect effect of the program could be to induce those with debts below the threshold to take action to remain below it, and not incur additional debts. Second, we note that program eligibility requires not just having debt above the threshold, but also that a taxpayer either has had a Notice of Federal Tax Lien filed and the associated Collection Due Process (CDP) hearing rights have expired, or a Notice of Levy has been issued. If taxpayers know that either of these two tools could trigger passport program eligibility and subsequent certification, taxpayers may respond more quickly to them. Either of these responses would increase tax collection.

2.10.1. Indirect effect related to the statutory threshold

To understand the indirect effects on initially ineligible taxpayers (those with debt below the thresholds), we study the behavior of taxpayers around the eligibility threshold. The hypothesis for this particular indirect effect is that, after the introduction of the passport revocation policy, the cost of having tax debt above the eligibility threshold increased, and thus those with tax debt below the threshold have an additional incentive to remain below that threshold. (Recall there is no reason to observe exceptional bunching from above, as reducing one's debt to just below the threshold does not lead to decertification.) We perform an initial test for this effect by looking to see if taxpayers "bunch" below the (changing) threshold. One important thing to note is that, for passport program eligibility, what matters is *assessed* balance, penalties, and interest. Although additional penalties and interest may be accruing over time, these *accrued* amounts do not affect program eligibility. That is, taxpayers could target their assessed balance, penalties, and interest due below the threshold, and maintain those balances for some time.

The data do not suggest that taxpayers are bunching under the passport thresholds. Figure 2.19 presents histograms of total assessed balance in December of each year 2010 through 2019, grouping by taxpayers in each $1,000 balance bucket. Although there is indeed evidence of bunching below $50,000, this pattern starts well before the passport program was passed into law in 2015, and certainly before its rollout and implementation after that. It is also notable that the bunching below $50,000 persists even when the passport threshold begins rising above that figure. Our hypothesis for this pattern of behavior is that the observed bunching is driven not by the passport program, but instead is related to the streamlined Installment Agreements and other Installment Agreement policies that use a $50,000 threshold that is constant over this period. The passport program could still affect this pattern, however; if the potential certification for passports makes payment plans more attractive, this could increase the incentive to remain below $50,000, and thus the amount of observed bunching there.

Figure 2.19: Histogram of taxpayers by total assessed balance

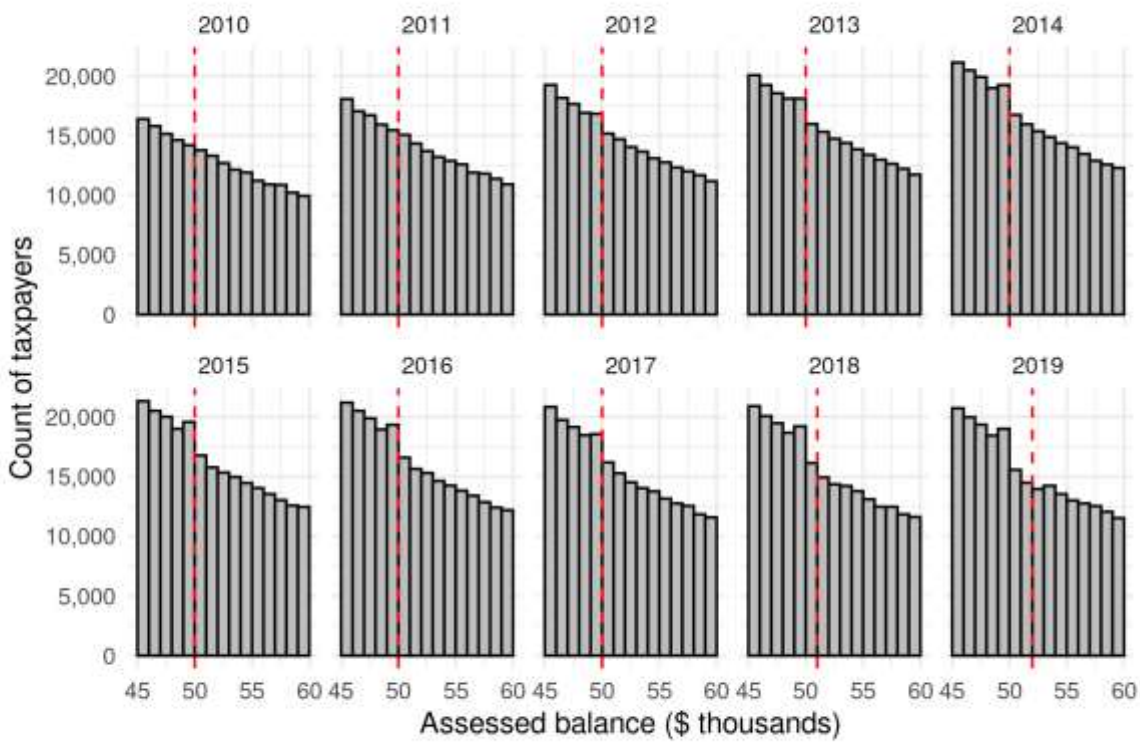

Notes: This figure plots the count of taxpayers in each $1,000 bucket by total assessed balance, in December of each year. The red lines depict the passport threshold: $50,000 in every year until rising to $51,000 in 2018, and $52,000 in 2019.

2.10.2. Indirect effect related to other enforcement tools

Another potential mechanism by which the passport program could have an indirect effect on taxpayer compliance is by increasing the efficacy of other existing enforcement tools, especially liens and levies. In addition to having a debt balance above the statutory threshold, a prerequisite for passport certification is that a taxpayer either has had a Notice of Federal Tax Lien filed and the associated Collection Due Process (CDP) hearing rights have expired, or a Notice of Levy has been issued. Each of these tools (liens and levies) has its own effect on taxpayer compliance (see, e.g., Turk et al. (2016) and Collins et al. (2018)). We hypothesize that the passport program makes these tools more effective, especially for taxpayers with a balance above the passport program's statutory threshold. For such a taxpayer, the imposition of a new lien or levy would also carry the cost associated with the potential for imminent passport certification, and thus might induce a stronger response than a lien or levy imposed prior to the passport program.

Credible identification of these more subtle indirect effects of the passport program would require overcoming significant statistical challenges, including the endogeneity of liens and levies and disentangling the various other policy and enforcement changes that occurred in the years around the announcement and implementation of the passport program. For the present study, we simply note the issue and suggest it as a promising avenue for future research.

2.11. Appendix C: Sample letters

Sample certification and decertification letters are available online:

508C (Certification): https://www.irs.gov/pub/notices/cp508c_english.pdf

508R (Decertification): https://www.irs.gov/pub/notices/cp508r_english.pdf

Chapter 3. Does Shaming Pay? Evaluating California's Top 500 Tax Delinquent Publication Program (with Chad Angaretis, Brian Galle, and Allen Prohofsky)[74]

3.1. Introduction

Given the limited resources available to tax enforcement authorities around the world, tax agencies are always in search of cost-effective methods for collecting revenue and assuring tax compliance. Many countries and U.S. states use public disclosure of tax debtors' personal information (sometimes known as "name and shame" lists) to encourage increased tax compliance and collection of outstanding tax debt. To date, little is known about whether these programs are effective; only two existing studies have shed light on these questions (Perez-Truglia and Troiano 2018, Dwenger and Treber 2019). We add further evidence to the existing literature as well as answering new questions by studying California's Top 500 delinquent disclosure program, using restricted-access administrative tax data covering five years of twice-yearly list publications.

Delinquent taxpayer disclosure schemes could play an important role in a compliance regime, even if disclosure is limited only to the largest debtors. Perceptions that the rich and powerful comply can have significant impact on tax morale and therefore compliance (Slemrod, Ur Rehman

[74] We thank Jeff Hoopes, Joel Slemrod, and participants at the 2020 National Tax Association annual conference, ComplianceNet 2021, and 2021 IIPF annual congress for many helpful comments and suggestions. We are especially grateful to Nghia-Nhan Duong and You Zhan for exceptional research assistance, and to Ken Kulhavy, Alaina Andrews, Jeff Geisler, Cesar Ramos, and the rest of the CART team at the Franchise Tax Board. Any views expressed in this paper are those of the authors and not official positions of the California Franchise Tax Board.

and Waseem 2020). Disclosure strategies aimed at large tax delinquents can potentially help to build this perception.

More generally, non-monetary sanctions are a relatively unexamined but theoretically useful tool for a tax authority (Blank 2014, Kuchumova 2018, Kuchumova 2021, Organ, et al. 2021). Typically, monetary penalties dominate alternatives because they result in transfers, rather than deadweight loss (Polinsky and Shavell 2000). Almost by definition, however, persistent tax debtors are relatively insensitive to monetary penalties: if the authority has not been able to initially collect the tax liability, it is unclear why it would be better able to collect any additional penalty. Optimal enforcement theory suggests that multiple enforcement instruments can be desirable when the target population has heterogeneous sensitivity to each instrument (Slemrod and Gillitzer 2013), and this can hold even for transferless instruments in some cases (Galle and Mungan 2021). Imprisonment, although practiced in the United States as a means of securing some public debts, may lose money on net through its high cost and negative impact on earnings. Disclosure appears to be a low-cost alternative, if it is effective, although as we discuss later, the total measured cost of the program should include not just administrative costs, but also the potentially significant disutility costs incurred by those whose information is published.

To gain traction on some of these questions, we study several components of the California "Top 500" disclosure program, a semi-annual internet posting of California's largest tax debtors.[75] We observe outstanding balance, payments, and other administrative outcomes for California taxpayers with outstanding tax debt of at least $100,000 (well below the threshold for Top 500 publication, which ranges from about $150,000 to $230,000 during our study period). In addition, we link these data to individual CA tax return information. This linkage allows us to condition

[75] The program includes taxpayers owing the top 500 liened delinquencies of personal income tax (PIT) or business entity (BE) tax in excess of $100,000. We study individual taxpayers.

responses on observed taxpayer characteristics, and also to measure the extent to which publication and other administrative steps aimed at collection of old tax debts affect reported income in subsequent tax years. Although California posts both individual and corporate debtors, we limit our focus to individuals.

Tax debtors in California receive several notifications before their names can be posted publicly. The "pre-letter," an initial notification about the existence of the Top 500 list and the possibility that any debtor who owes in excess of $100,000 may appear on it, is sent semi-annually to all $100,000+ debt households who do not fall into a statutory exception, usually arriving shortly after the publication of the most recent list. Several months later, and two months before posting the final list, a second letter (referred to by staff as the "official letter") is sent to the 500 publication-eligible taxpayers with the highest debts outstanding at that time, informing these taxpayers that their information will be published if they do not take action. Finally, there is the published list itself, which provides names, addresses, unpaid balance amount, and professional license information for all taxpayers who received the official letter and did not take sufficient action to avoid publication.

Our main analysis exploits the random variation in letter cutoffs to estimate the impact of receiving the letter, controlling for debt levels. Over the 10 list publication cycles in our study period, the lowest balance we observe receiving an official letter each cycle ranges from roughly $150,000 to $230,000. This lowest balance is determined by the 500th-largest eligible balance, which varies with each cycle. In this specification, identification is based on observing two taxpayers with similar balance, one who is mailed a letter and the other who is not. Because the cutoff is stochastic, even taxpayers who are aware of the program cannot perfectly predict whether they will receive a letter, reducing any concern about selection into treatment status.

We take these identification methods to a set of compliance outcomes. Most simply, we look at the extent to which treated taxpayers make any payment, or pay more or less than others. We also examine other behavior of interest to tax administrators, namely entering into an Installment Agreement to pay down the total balance over time, or taking steps necessary to qualify for other statuses leading to exemption from publication, such as filing for bankruptcy or documenting other significant financial hardship.

We find that taxpayers are responsive to receiving an official letter. Our preferred identification strategy suggests that, over the two-year period following treatment, recipients pay an average of about $7,200 more than non-recipients, yielding about $2.8 million in added revenues for each year we observed the program in operation. Because this specification is restricted to observations with balances in the range of cutoff values (roughly $150,000 to $230,000), this number omits the highest-balance households. Estimates using our full sample suggest total added revenues of $7.2 million annually. Drawing on inferences from the behavior of marginal avoiders, we attempt to put bounds on the private costs of the program, including the welfare costs of being subject to publication. Combining these with our revenue estimates allows us to further estimate that this revenue results in net social welfare gains of about $1 million per year.

A large fraction of treated households also take steps to make themselves ineligible for publication, with an increase in ineligibility of eight percentage points relative to the control households in the three months after treatment. Over the two-year window, the increase rises to twenty percentage points. This is on a baseline that by construction has zero ineligible households. Of these, two percentage points enter into Installment Agreements in the three months after treatment, and over two years there is a 12 percentage point rise in such agreements.

Our paper builds on the two existing studies about delinquent taxpayer disclosure. Perez-Truglia and Troiano 2018 (henceforth "PTT") used randomized letters sent by the researchers to highlight the salience of published tax delinquents' information being public. Letters increased the probability that low-balance tax debtors (those with debts below $2,500) would leave the list but had no effect on high-balance tax debtors.[76] Dwenger and Treber 2019 (henceforth "DT") study the first year of a delinquent taxpayer disclosure program for corporations and self-employed individuals in Slovenia. Studying responses after the program was announced but before it was implemented, DT find that the threat of publication leads both corporations and the self-employed to reduce tax debt. Publication itself led to further reductions in tax debt, though of a smaller size.

We make several new contributions. First, by combining payments data with individual income tax returns, we observe whether the treatment has any effects on (reported) subsequent income, as well as conditioning other responses on reported income, self-employment, and tax filing status. This allows us to begin to untangle *why* the studied households fail to make timely payments. The self-employed are much more responsive to treatment, suggesting that failure to pay assessed taxes is tax avoidance, rather than the result of budget pressure, for those not subject to withholding. We also find a significant and economically substantial positive correlation between failure to file any return and non-compliance; this could reflect a high subjective cost of compliance, perhaps due to filing complexity, or that some households may simply have a relatively strong preference for avoiding contact with California tax authorities (prior work has found that complexity is an important contributor to non-filing behavior (Bhargava and Manoli 2015)). Ability to pay does play a role for some households, however. We find that the effect of

[76] Other field studies find that mailings can increase compliance by raising the salience of detection or sanction (Bott, et al. 2019, Cranor, Goldin and Homonoff 2020, Dorrenberg and Schmitz 2017, Gemmell and Ratto 2018, Holz, et al. 2020, Iyer, Reckers and Sanders 2010).

the official letter on most of the compliance actions we study, including the total payment triggered by treatment, is considerably higher among households reporting the highest adjusted gross income.

In addition, we provide evidence about a program already in place, studying the effect of the program doing "business-as-usual," in contrast to PTT's study of the effect of outside researchers augmenting a program's salience, or DT's study of the rollout of a new program, when public attention on the program may be abnormally high. Third, we offer evidence on a program targeting only high-balance individual tax debtors, in contrast to PTT's study of three states using relatively low balance thresholds for publication (publishing all tax debtors with debts above $250, $2,500, or $5000, depending on the state) or DT's study of only self-employed individuals. Lastly, our study includes analysis of payment and specific non-payment compliance outcomes, in contrast to PTT's aggregate compliance outcome of no longer appearing on published lists, which cannot distinguish between the various reasons an individual might be removed from the lists.

The rest of the paper proceeds as follows. Section 3.2 describes California's Top 500 program in further detail, and Section 3.3 describes the data available for this study. Section 3.4 is the main section of the paper, describing our analysis of the effects of the official letter notifying taxpayers of their imminent publication. Section 3.5 then describes a brief analysis of the effects of publication and license suspension. In Section 3.6 we set out a theoretical framework for evaluating the outcomes we measure, and in Section 3.7 apply that framework. Section 3.8 concludes.

3.2. Overview of the Top 500 program

California imposes a progressive income tax with a top rate of 13.3%. The state generally follows federal rules for defining the tax base, with certain exceptions. The state's income tax is administered by the California Franchise Tax Board ("FTB"). One tool FTB uses for collecting unpaid tax liabilities is its "Top 500" program. Legislation enacted in 2006 (AB 1418) mandated that FTB annually make public a list of its 250 largest debtors. AB 1424, enacted in 2011, expanded the list to the top 500 debtors, increased the frequency to twice a year, and added additional sanctions for listed debtors, including provisions for the suspension of professional licenses and a prohibition on being awarded state contracts.[77]

California begins assembly of its Top 500 list with a preliminary list of all taxpayers with current unpaid balances of more than $100,000, a group that typically numbers about 6,000 households. FTB staff then scrutinize this preliminary list in an effort to identify taxpayers who are statutorily exempt from being included in the Top 500. The most common exempt categories are for taxpayers who have entered into a payment agreement with FTB or been found to suffer from financial hardship.[78] Others include deceased individuals, "innocent spouses" not responsible for the household's debts, and taxpayers who have commenced federal bankruptcy proceedings.

The winnowing process typically leaves approximately 3,000 eligible individual taxpayers. We call this group the "pre-letter list." At this point FTB prepares a mailing list. Taxpayers who remain on the pre-letter list and have never previously been included in the Top 500 receive a letter (Appendix, Figure 3.13) informing them of their status as a potential Top 500 includee.

In general, taxpayers with unpaid balances of this magnitude have already been the subject of extensive collections efforts. Debts are only counted towards the delinquent total if they have been

[77] The two pieces of legislation created parallel, but separate, programs for California's top income tax debtors and its top sales tax debtors. California's Sales and Use Tax is administered by a different agency; we do not have access to their data. The change from publishing the top 250 to the top 500 is also of research interest, but because our data begin in 2013 we are unable to study that in this paper.

[78] FTB defines financial hardship as net assets that are less than necessary to "provide for the [taxpayer's] health and welfare" or the "reasonable costs … of the [taxpayer's] trade or business." California Code of Regulations § 19195-2. In addition, the relevant statute requires FTB to remove a taxpayer from the list promptly when it determines that the delinquency is "uncollectible." CA Rev. & Tax Code § 7063(f)(4). FTB cannot collect debts more than 20 years old, with certain exceptions. CA Rev. & Tax Code § 19255.

delinquent more than 90 days, and the state has filed a notice of tax lien. All taxpayers with balances above $25,000 and who have failed to reach compliance voluntarily are assigned to an individual agent at FTB, who attempts to contact the taxpayer and work with them personally to collect the outstanding obligation. Taxpayers who do not comply at this point are also subject to wage garnishment. This accounts for the relatively large share of accounts with high balances being excluded from the pre-letter list due to payment agreements, hardship findings, or other ineligibility.

After sending the pre-letter, FTB staff then begin a more thorough review of the potential set of taxpayers who will become the Top 500. As more information is gathered about taxpayers, some of those who received pre-letters are subsequently deemed ineligible for publication. Taxpayers also may take actions after receiving a pre-letter which lead to ineligibility for publication, including entering into payment agreements. Following this review and action, taxpayers who are still eligible for publication are ranked from highest to lowest balance, and a second mailing list is prepared.

The 500 taxpayers with highest outstanding balances are then sent a letter (the "official letter") informing them that if full payment or other resolution is not made in the next 60 days, they will be included in the Top 500 list posted online. Because the Top 500 is based on all balances due (among both individuals and businesses) the number of individuals sent this letter is less than 500; typically, about 400 individuals receive the official letter. The letter provides taxpayers with information about how to contact FTB to resolve their tax debt, and FTB's web site provides a link taxpayers can click to initiate contact.

Sixty days following the official letter, after a final check to confirm all individuals remaining are eligible for publication and have not taken actions that would lead to ineligibility, FTB posts the remaining taxpayers' information online. FTB does not replace taxpayers who become ineligible between the time of the letter and the list, so that taxpayers that pay their balances or enter into payment agreements during this time reduce the size of the list. In practice, most posted lists include about 300 individuals, providing some initial evidence that the official letter is fairly effective.

This process is carried out twice a year. Figure 3.1 summarizes the Top 500 timeline. Immediately after the Top 500 List posts, the process repeats with another gathering of taxpayers who then have $100,000 or more in debt.

Figure 3.1: Top 500 timeline

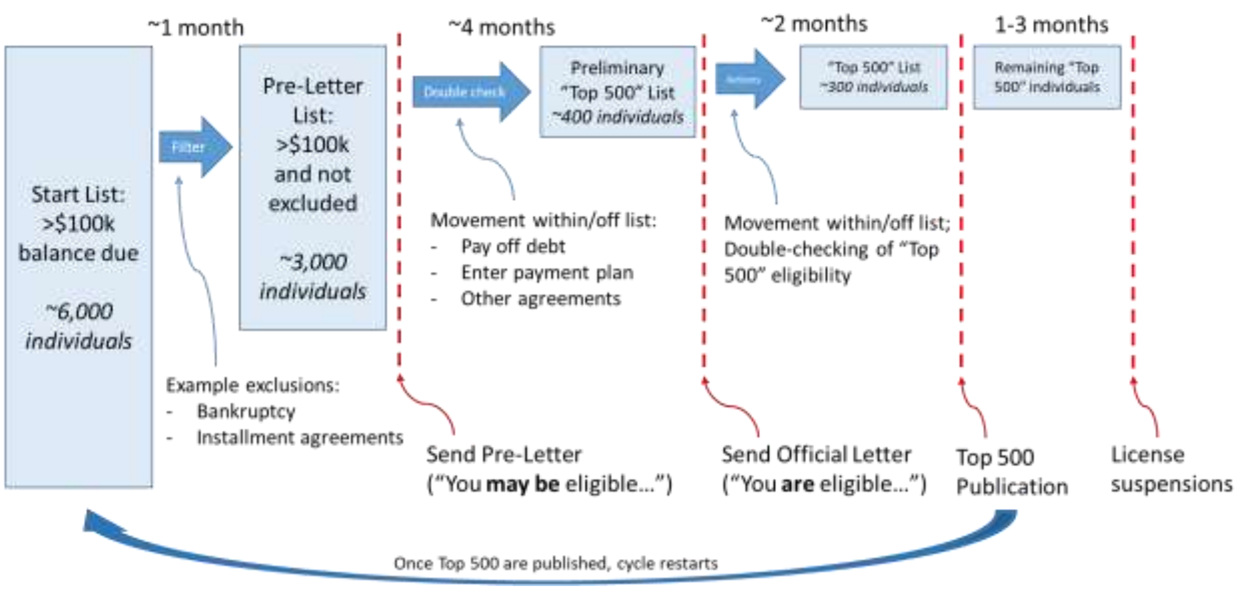

Notes: This figure describes the timeline of a typical Top 500 publication cycle, from start to finish. The Top 500 list is published twice per year, in April and October.

Appearing on the list triggers additional penalties on top of public disclosure of the taxpayer's name and debt. In most cases listing triggers suspension of professional, occupational, and even driver's licenses. Most professional associations and licensing agencies cooperate with the FTB to suspend licenses, with one notable exception. A license to practice law is not automatically suspended, but the State Bar of California may recommend suspension at its discretion (CA

Business & Professions Code § 494.5). In practice, the State Bar does not suspend licenses to practice law for nonpayment of taxes. State agencies cannot enter into contracts with taxpayers who appear on the list.

There is an additional process that licensing entities must follow before suspending a license. The licensee must receive a separate notice of license suspension within thirty days of appearing on the list. The licensee can obtain a temporary license (if in the application period) for ninety days, but at between ninety and one hundred twenty days after mailing of the suspension notice, the license is suspended. For individuals who are published on the Top 500 list, we observe whether a licensing entity has notified FTB that a licensed individual is on a list, but we do not have additional information on whether the licensing entity complies with the additional notice procedures.

3.3. Data

Our data comprise a merged set of payment, balance due, and other individual-level tax information for every California taxpayer who has incurred a balance due of at least $100,000. We observe each time a taxpayer appears on FTB's initial list, receives a pre-letter, receives an official letter, or makes the Top 500 list. We also observe all payments made, as well as status and activity code data that allow us to observe other outcomes of interest, such as entering into an installment agreement.

We can also match these payment and status records with tax filing information for each taxpayer who ever appeared on the initial FTB lists (i.e., those who at one point had a delinquency of at least $100,000). This provides us with selected fields on the taxpayer's California (not federal) individual income tax returns, stretching from 2009 to 2019; however, about forty percent of the household-years we observe in the payments data lack a timely return for the two years before observation.[79]

Summary statistics for the individuals receiving the official letter, across all 10 cycles we study, are shown in Table 3.1.[80] The mean balance among letter recipients is about $859,000, while the median is lower at about $324,000. When restricting to first-time letter recipients only, we see slightly lower balances, with a mean of $606,000 and median of $300,000. Slightly more than half of letter recipients have filed returns for the years just prior to their letter. Among those filers, there is a wide range of income. About half of those with filings report some business income. Nearly all of the letter recipients are California residents.

Table 3.1: Summary statistics for official letter recipients

	Mean	Std. Dev.	P5	P25	Median	P75	P95
Panel A: Among all official letter recipients							
Balance due as of official letter ($ thousands)	859	7,069	181	245	324	539	1,796
Filed on-time return for two-years prior tax year (1/0)	0.57	0.50	0.00	0.00	1.00	1.00	1.00
Filed return for prior tax year (1/0)	0.55	0.50	0.00	0.00	1.00	1.00	1.00
Among those with filed returns for prior tax year:							
AGI ($ thousands)	-250	3,913	-1,100	2	40	152	884
Wages ($ thousands)	71	699	0	0	0	36	213
Has business income (1/0)	0.49	0.50	0.00	0.00	0.00	1.00	1.00
CA resident (1/0)	0.96	0.19	1.00	1.00	1.00	1.00	1.00
Panel B: Among first-time official letter recipients							
Balance due as of official letter ($ thousands)	606	2,012	172	229	300	493	1,512
Filed on-time return for two-years prior tax year (1/0)	0.56	0.50	0.00	0.00	1.00	1.00	1.00
Filed return for prior tax year (1/0)	0.53	0.50	0.00	0.00	1.00	1.00	1.00
Among those with filed returns for prior tax year:							
AGI ($ thousands)	-111	3,509	-825	2	41	165	983
Wages ($ thousands)	98	943	0	0	0	43	272
Has business income (1/0)	0.49	0.50	0.00	0.00	0.00	1.00	1.00
CA resident (1/0)	0.97	0.17	1.00	1.00	1.00	1.00	1.00

[79] A taxpayer who fails to file any return may still show a balance owed to California as a result of information return reporting from third parties. Non-filers are likely over-represented among the group of major tax delinquents because they fail to report any potential offsetting deductions, inflating their tax due.
[80] To maintain anonymity with respect to data that are not disclosable under the Top 500 statute, we report descriptive and other statistics in bins large enough to prevent individual identification.

Notes: This table presents summary statistics for official letter recipients, across all 10 cycles included in our study.

For each publication cycle, we observe the lowest balance among the letter recipients for the cycle, and we call this the "cutoff" value for that cycle. Cutoffs for the ten cycles we observe range between \$150,000 and \$230,000, as illustrated in Figure 3.2.

Figure 3.2: Variation in official letter cutoffs over time

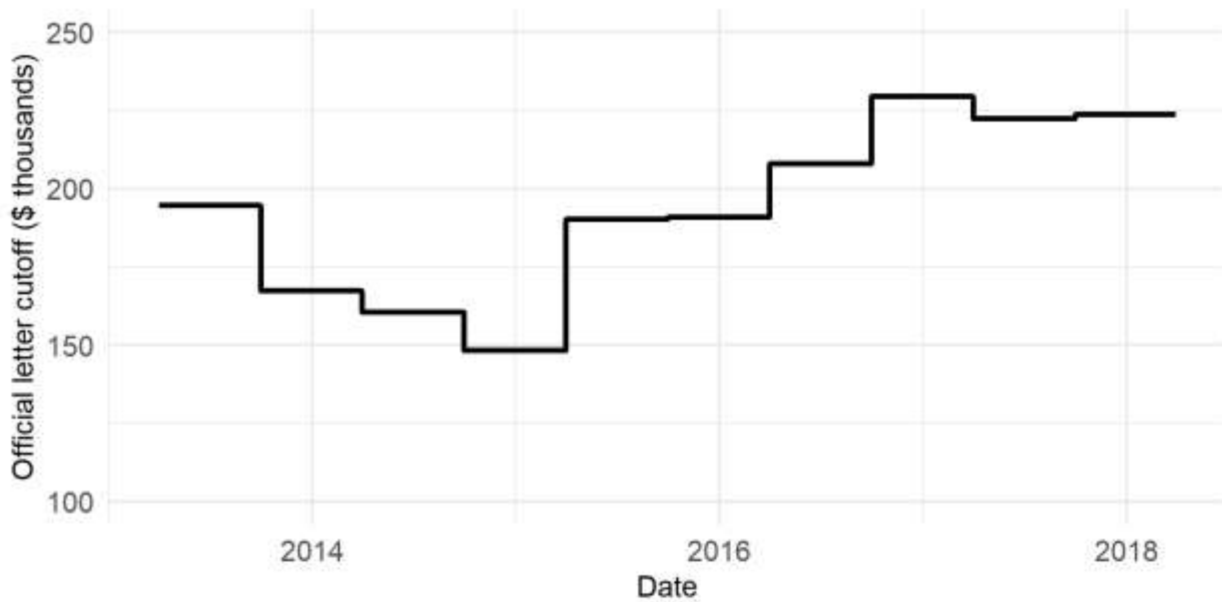

Notes: This figure plots the lowest observed balance among official letter recipients each cycle. Because the official letter is sent to the limited number of taxpayers likely to be in the Top 500 balances, and the set of tax debtors changes over time, the cutoff for official letter receipt also changes over time, as shown here.

In our reported results, we screen out households we can identify as statutorily ineligible for publication. By definition, our "treated" households have been found eligible. To help ensure comparability between treated and control households, we attempt to similarly limit the control population to those who would be eligible if their balance were high enough. Our data do not directly report eligibility. We do, however, observe the "status codes" that FTB uses to determine eligibility, such as whether a taxpayer has entered into an installment agreement, a bankruptcy, or has established that they are an innocent spouse. We thus omit households with one or more of these observed status codes prior to treatment. After screening, we still observe some taxpayers with balances above the cutoff who do not receive letters, suggesting that our screen does not map perfectly onto ineligibility.

A subset of taxpayers remain on the Top 500 list persistently. On average, a little under half of those receiving the official letter received one for a prior cycle, and again slightly less than half of those ultimately published each cycle were published on a prior cycle as well. This pattern can be seen below in Figure 3.3. Households already published in a prior cycle may be especially unlikely to respond to treatment, as by definition they have already failed to do so once before. Including them in our control (treatment) population might therefore bias our measured results upwards (downwards). Unless otherwise noted, our reported results therefore omit households published in a prior cycle.

Figure 3.3: Official letter and publication counts, initial vs. repeat appearance

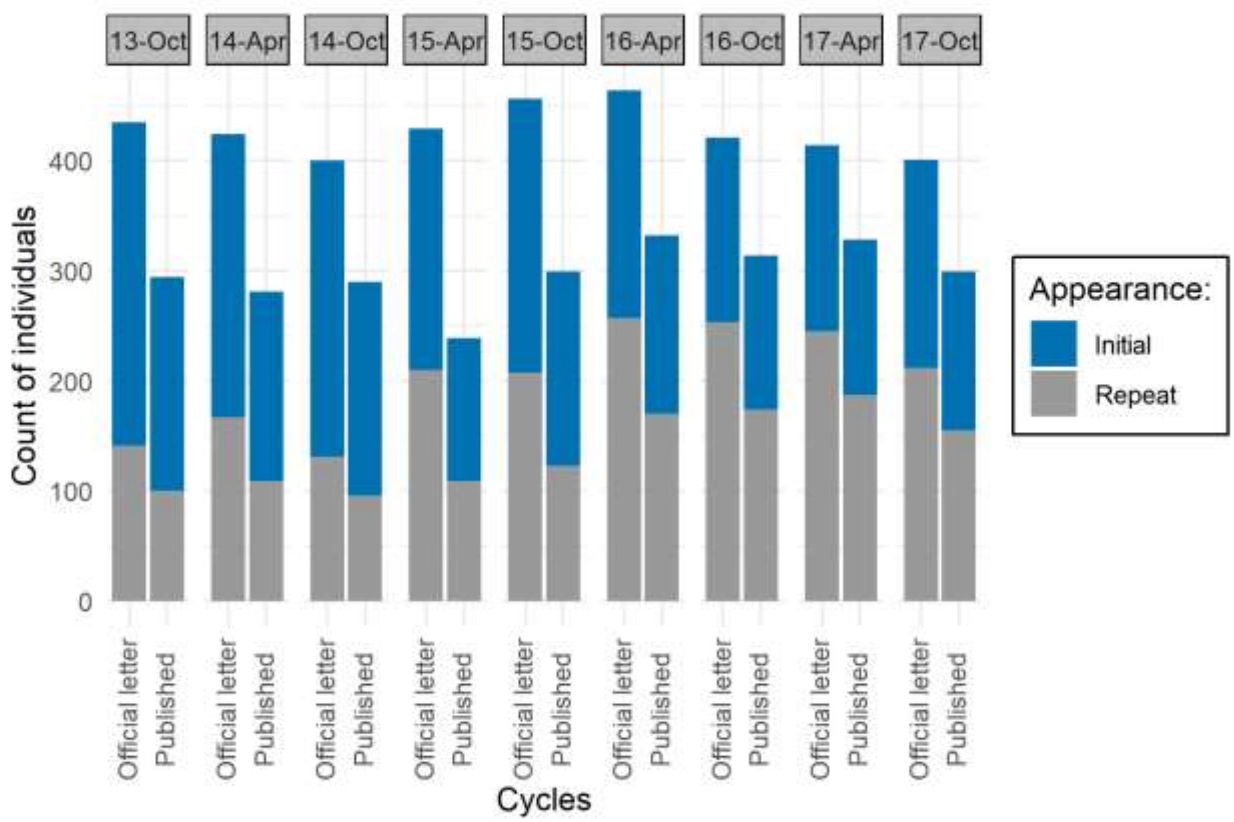

Notes: This figure plots the count of individuals receiving the official letter, and the number ultimately published, each cycle from Oct. 2013 to Oct. 2017. Individuals are counted separately by whether or not it is their first time receiving an official letter or getting published.

3.4. Effects of the official letter

The central treatment we study in this paper is the official letter, by which taxpayers are notified that they are slated to appear on the Top 500 list. This notice is sent four months after the pre-letter mailing, and two months before list publication, and is sent only to the 500 taxpayers with the highest balances among those eligible to appear on the list. The letter is a credible and time-sensitive notification that a taxpayer will be published if they do not act, and so we expect it to have the largest impact. Because our data include codes for a bad mailing address, and we omit these observations, we can verify that we are measuring true treatment effects and not simply intent to treat.[81]

We note that the sample available for analysis of the official letter treatment has already been selected on their being non-responsive; by definition, these are taxpayers who have failed to respond to a series of prior notifications, including a notice of tax lien. In addition, all of these households have received a "pre-letter" notifying them that they may qualify for the Top 500, giving those that are most responsive an opportunity to take action prior to the official letter.[82] Our

[81] Typically, fewer than 10% of the individuals on the "start list" of individuals considered for the pre-letter have such bad address codes. This represents a much lower share than in other reported tax authority mailing interventions (e.g., Goldin, Homonoff, et al. 2021), consistent with our treated population already having been the subject of close human interaction. Another concern with letter studies is that some taxpayers may fail to open or read the letter (e.g., Perez-Truglia and Cruces 2017, Bottan and Perez-Truglia 2020, Nathan, Perez-Truglia and Zentner 2021). To the extent this is true, our estimates reflect a lower bound on the treatment effect. Our setting is somewhat distinctive from pure mailing interventions, however, in that letter recipients who fail to read the letter are treated later through a non-mail treatment, namely publication. As described below, we find little incremental impact of publication, suggesting that most of those who are susceptible to treatment are reading the letter.

[82] Certain data on individuals with balances between $75,000 and $100,000 are available to the researchers. A preliminary analysis using these data to study the effect of the pre-letter suggests it does not have a substantial effect. Difference-in-discontinuities analysis around the $100,000 threshold for pre-letter receipt cannot reject zero effect for most outcomes, except that for April pre-letter recipients, payment amounts may be marginally higher.

results in this section can thus be understood as a lower bound on the effect of disclosure on populations that have not been as thoroughly pre-selected.

3.4.1. Graphical analysis

We start with a simple graphical analysis showing behavior before and after the official letter dates, for letter recipients and non-recipients. In this section we focus our attention on the taxpayers who are eligible for publication on the Top 500 list and who have not yet received an official letter (that is, for taxpayers who remain on the list for multiple cycles, only their first observation is included). This allows us to present a visual test of the effect on a taxpayer of encountering, for the first time, the letter and its notification that the taxpayer will be published if action is not taken quickly.

Figure 3.4 below compares official letter recipients and non-recipients on four behaviors: making a payment, average payment amount, starting an Installment Agreement, and entering into a status that makes one ineligible for Top 500 publication. In this figure, we restrict to those within the cutoff range (roughly $150,000 to $230,000), where (as we describe more below) the argument for quasi-random treatment is strongest.

Figure 3.4: Behavior before and after the official letter, recipients vs. non-recipients

(only those within cutoff range)

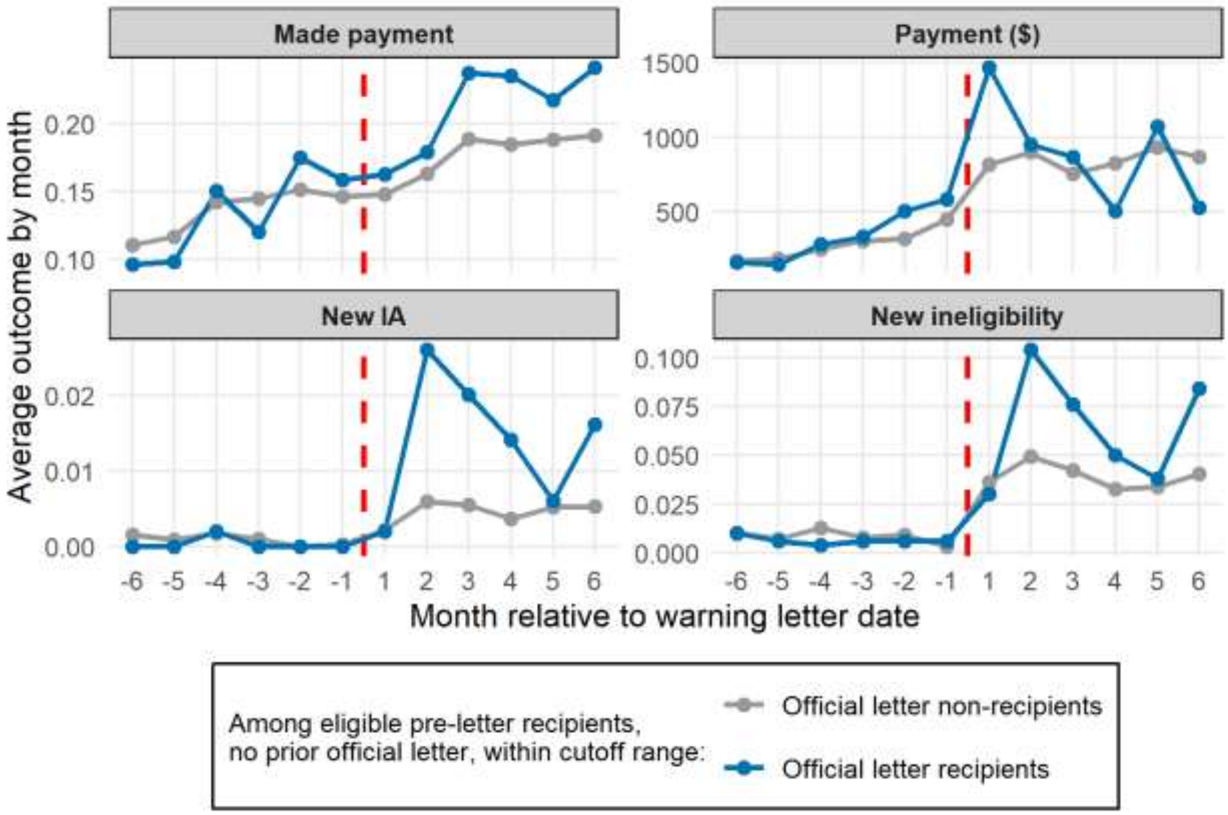

Notes: This figure compares the behavior of two groups of taxpayers around the date of the official letter. In gray are taxpayers who are eligible for publication and received a pre-letter, but did not receive an official letter. In blue are similar individuals (eligible pre-letter recipients) who did receive an official letter. In this figure we exclude individuals who previously received an official letter (i.e., we focus on first-time recipients), and we limit to those with balances between $150,000 and $230,000, the range of cutoff values inside of which treatment is quasi-random.

The figure suggests that the official letter has a strong effect. Average payment amounts jump after treatment among treated households. There is also a sharp rise in the share of treated taxpayers entering into new Installment Agreements, and more generally taking actions leading to ineligibility for publication. As for the share making payments, the above-trend but delayed effect (the gap between treated and untreated widening after three months) could be related to the Installment Agreements taking time to set up and first payments to begin.

The small break in trend after the letter for the control group is a mechanical effect: inclusion in the control group, like inclusion in the treatment group, is conditional on eligibility for the letter, which necessarily means no compliance actions have been taken in the preceding months. Conditioning on an action not happening in the past means that the probability of it occurring in the next period is likely to jump; the patterns here show this, with the effect of the letter demonstrated by the additional increase in activity for the treated group above and beyond the mechanical effect for the control group.

We find similar, and in fact stronger, effects when examining the full sample (see Appendix Figure 3.17). Seeing stronger effects relative to the effects among those in the cutoff range suggests that the official letter has larger effects on those with higher balances. Because we cannot compare treated high-balance debtors to untreated taxpayers with similar balances, however, we cannot fully rule out the possibility that full-sample results are caused by some unobserved phenomenon that happens to affect only high-balance debtors at just the time of treatment.

We also test the sensitivity of these patterns to our eligibility and other data filters. Appendix Figure 3.18 shows that the patterns are consistent when removing the pre-letter, initial receipt, and eligibility filters that are applied in our main specification above.

3.4.2. Random cutoff analysis

To develop a more precise, quantitative estimate of the effect of the official letter, we now turn to a regression analysis that exploits the variation in the official letter cutoff. Across the ten cycles in our data, the lowest balance receiving an official letter ranges from approximately \$150,000 to \$230,000, as shown in Figure 3.2. Because the cutoff is determined by the 500th highest eligible balance, the cutoff dollar value cannot be predicted precisely in advance. As individuals are accruing and paying down balances over time, independently, influenced by myriad factors unrelated to the publication program (for example, volatile income, liquidity constraints, and fluctuating asset values, to name a few), the ranking of balances and value of the 500th highest eligible balance changes such that individuals who are close to the range of historic cutoff values cannot know for sure whether they will be on the list or not. A publication-eligible taxpayer with $175,000 of balance as of the official letter date in one cycle might receive the letter, while in another cycle a taxpayer with the same balance would not. In effect, taxpayers randomly assigned to a cycle in which they do not receive a letter serve as controls for taxpayers with the same balance who are randomly assigned to a cycle in which that balance does trigger a letter. This mitigates the possibility of selection into treatment.

We use this quasi-random variation to estimate the effect of official letter receipt. We start with a pooled difference-in-differences approach, as follows:

$$Outcome_{itc} = \alpha + \beta_1 \cdot Post_{itc} + \beta_2 \cdot Treat_{itc} + \beta_3 \cdot Post * Treat_{itc} + \gamma \cdot Balance_{ic} + April_c + \varepsilon_{itc}$$

In this specification, i indexes individuals, *c* the cycle in which we observe them, and t indexes the month relative to the official letter mailing date, ranging from -6 to 6. $Post$=1 for months after the letter, $Treat$=1 for individuals receiving the letter, and $Post * Treat$ is their interaction. *Balance* does not vary over a given cycle and is measured at the date the official letter list is determined.[83] We also control for whether the observation is for an April or October Top 500 list cycle (to address the potential for seasonality in payments and other actions). We run this specification including only those taxpayers who have balances in the range of the cutoffs and who have never received an official letter before.[84] We test several outcomes, including three monthly binary variables: (1) starting a new Installment Agreement, (2) entering into any status that makes

[83] Because we observe compliant households only in one cycle, we omit individual-unit fixed effects.
[84] We conduct a variety of robustness analysis to confirm that the results are not sensitive to our data filters, regression specifications, or other choices. Appendix C includes these results.

one ineligible for publication, and (3) making a positive payment. We also test the non-binary outcome of monthly payment amount, in dollars.

Table 3.2 below reports the results of this specification for the period between zero and three months after treatment. In these initial estimates, we limit the sample to observations with balances between $150,000 and $230,000, so that every treated unit has at least one untreated control with a similar balance. The strongest measurable effects are on new installment agreements, and new ineligibility status more generally. By construction, there are zero households in these statuses at the time of treatment. During the three months after treatment, the official letter led an average of an additional 1.2 percent of treated households to enter new Installment Agreements each month, relative to untreated households. For the more general outcome of new ineligibility, the effect is an average additional 2.8 percent each month. Although we observe a positive coefficient, we do not find a statistically significant effect on payment amounts. Our point estimate for the probability of making a payment is positive but not significant at traditional levels.

Table 3.2: Official letter difference-in-difference results, observations within cutoff range

	Dependent variables:			
	Made payment	Payment amount ($)	New IA	New ineligibility
Official letter * Post	**0.0223***	**155.77**	**0.0119*****	**0.0283*****
	(0.0118)	**305.11**	**(0.0032)**	**(0.0066)**
Official letter	0.0010	73.73	-0.0004	-0.0004
	(0.0146)	173.93	(0.0003)	(0.0022)
Post	0.0192***	467.612***	0.0041***	0.036***
	(0.0028)	77.52	(0.0005)	(0.0016)
Balance	-0.0002	0.93	0.0000	0.0000
($ thousands)	(0.0002)	2.03	(0.0000)	(0.0000)
April publication	-0.0441***	-253.07***	-0.0014***	-0.0001
	(0.0049)	83.20	(0.0005)	(0.0015)
Intercept	0.1974***	319.08	0.0030	0.0099
	(0.0440)	356.70	(0.0022)	(0.0066)
Observations	37,848	37,848	37,848	37,848
R2	0.0047	0.0014	0.0041	0.0157
Mean dep var.	0.1583	606.2506	0.0029	0.0258

Notes: *p<0.1; **p<0.05; ***p<0.01. Standard errors clustered by taxpayer are shown in parentheses. This table reports the regression results for the main difference-in-difference specification. The underlying data include only those taxpayers eligible for publication who received a pre-letter for a given cycle and have balance within the range of cutoffs (roughly $150,000 to $230,000). The dollar value cutoff for official letter receipt depends on the 500th highest balance for a given cycle, and this varies across the ten cycles in our data as shown above in Figure 3.2. We thus argue that, within the range of cutoffs observed in our data, letter receipt is random.

We obtain similar results using an event study design, as follows:

$$Outcome_{itc} = \alpha + \sum_{t=-6}^{t=6} \beta_t Treat_{itc} + \gamma \cdot Balance_{ic} + April_c + \varepsilon_{itc}$$

The results from this specification, run on the same population as for the diff-in-diff above, are shown graphically in Figure 3.5 below, with full results presented in Appendix Table 3.5.

Figure 3.5: Official letter event study treatment coefficients

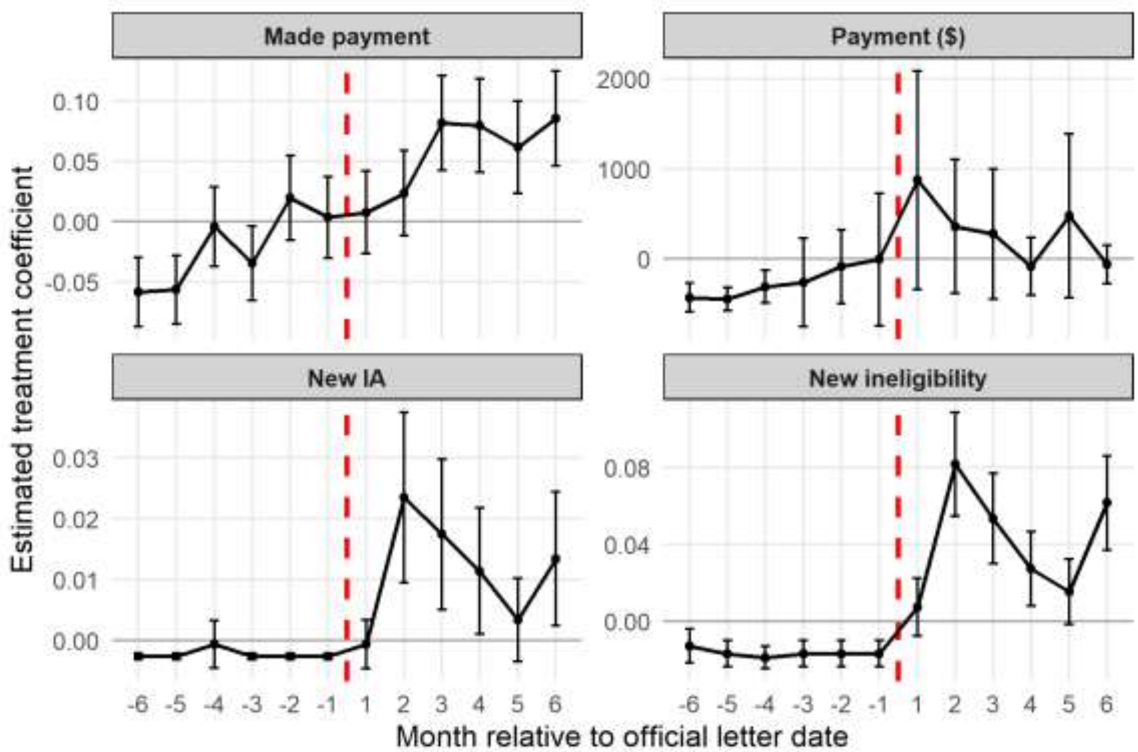

Notes: This figure plots the estimated coefficients on official letter receipt (treatment) dummies by month, from the event study regression approach.

3.4.3. Full sample analysis

A difficulty with the design we have pursued so far is that it does not allow us to test the effects of treatment for balances above the highest historic cutoff. Because balances are rightward skewed, the bulk of the unpaid debt lies in this region, and so responses by these households are of considerable policy interest. We therefore repeat our analysis with these taxpayers included. In this set of specifications, we cannot ensure that each treated unit has a matched control, but we can at least control parametrically for balance.

Because treatment is not randomly assigned for the upper tail of the distribution, we are more likely to have selection into treatment. That is, taxpayers with high balances are treated because they chose not to comply between the pre-letter and official letter. If anything, though, this selection effect reinforces our result. Those who do choose to be treated are those who are less apt to comply. We therefore are estimating the effects of treatment on a relatively unresponsive subset of the population. We cannot rule out the possibility, though, that some unobserved confounding event affects only high-balance treated households around the same time as the official letter.[85]

Table 3.3: Official letter difference-in-difference results, full range of observations

	Dependent variables:			
	Made payment	Payment amount ($)	New IA	New ineligibility
Official letter * Post	**0.0392*** (0.0063)**	**1621.5*** (471.06)**	**0.013*** (0.0018)**	**0.0271*** (0.0035)**
Official letter	-0.05*** (0.0073)	(220.47) (216.96)	-0.0005*** (0.0001)	0.0007 (0.0012)
Post	0.0193*** (0.0016)	590.31*** (82.80)	0.004*** (0.0003)	0.0379*** (0.0009)
Balance ($ thousands)	0*** 0.0000	0.2771 0.1811	0*** 0.0000	0.0000 0.0000
April publication	-0.0414*** (0.0027)	-497.45*** (99.07)	-0.0007** (0.0003)	-0.0001 (0.0009)
Intercept	0.196*** (0.0047)	629.87*** (69.95)	0.001*** (0.0002)	0.0066*** (0.0006)
Observations	126,444	126,444	126,444	126,444
R2	0.0048	0.0018	0.0042	0.0167
Mean dep var.	0.1803	842.2733	0.0031	0.0268

Notes: *p<0.1; **p<0.05; ***p<0.01. Standard errors clustered by taxpayer are shown in parentheses. This table reports the regression results for the main difference-in-difference specification. The underlying data include only those taxpayers eligible for publication who received a pre-letter for a given cycle, but without the restriction on balance from the main specification. This allows us to now include the higher balances that make up most of the Top 500 list and are of greater policy interest.

The effects on average payment amounts in the three months after treatment are much larger in the full sample population, with a 95% confidence interval for the mean monthly treatment

[85] When we estimate including households that are above the cutoff but ineligible for publication, we obtain essentially the same results. In this specification, high-balance ineligible households serve as an additional control for the high-balance treated households. To be sure, there are reasons to believe that ineligible and eligible households would respond differently to treatment. But what we can say is that any unobserved confounder that is driving our results would have to affect only those high-balance households that are treated, and do so at around the time of treatment.

effect that runs from about $700 to $2,500. New Installment Agreements and new ineligibility determinations are very close to the restricted-sample estimates.

3.4.4. Heterogeneous effects

It is also of interest to explore whether treatment effects vary based on observable taxpayer characteristics. For example, as Kuchumova (2021) argues, disclosure is more likely efficient if it disproportionately affects high-earning households, as we expect based on the results in DT. We similarly expect to see larger results for filers with business income: businesses are likely more subject to reputational pressure, and non-business filers are more likely to have been subject to withholding or wage garnishment, leaving less room for them to change behavior in response in treatment. Holding these other factors equal, taxpayers who have already exhibited a relatively high subjective cost of compliance, such as by failing to file any tax returns at all, are also likely to be less responsive.

We therefore re-estimate our regressions from prior sections, this time conditioning on three key data points from our linked individual income-tax data: income levels (using CA AGI), the presence of business income, and whether the household filed a return for the prior tax year. For AGI and business income, we use values from the tax return filed in the same year as treatment (and thus exogenous to treatment, reflecting actions from the year prior). Thus, for the April and October 2015 cycles, we use income reported for the 2014 tax year, and we record the household as having filed on-time if they filed a return in 2014 for the 2013 tax year.[86]

We first summarize results graphically, in Figure 3.6. For each of our four main outcomes, the figure plots results by sub-group: balances above and below $250,000; on-time filers and non-

[86] We use filing for the prior year as a marker of on-time filing because we cannot observe exactly when in a year a return is filed. Thus, for April 2015, we would not be able to tell if a return for the 2014 tax year was filed in April 2015 or instead in December 2015.

filers; AGI above and below median (roughly $40,000); and whether the taxpayer reports any business income. We see that the average-payment response is strongly correlated with all of these groupings: high balance, high AGI, on-time filing, and business income all appear to correlate with higher post-treatment payments. Similar differences emerge for installment agreements and ineligibility status generally, though balance does not seem to matter for these outcomes. Business income also appears uncorrelated with ineligibility.

Figure 3.6: Behavior around official letter, splitting on tax characteristics

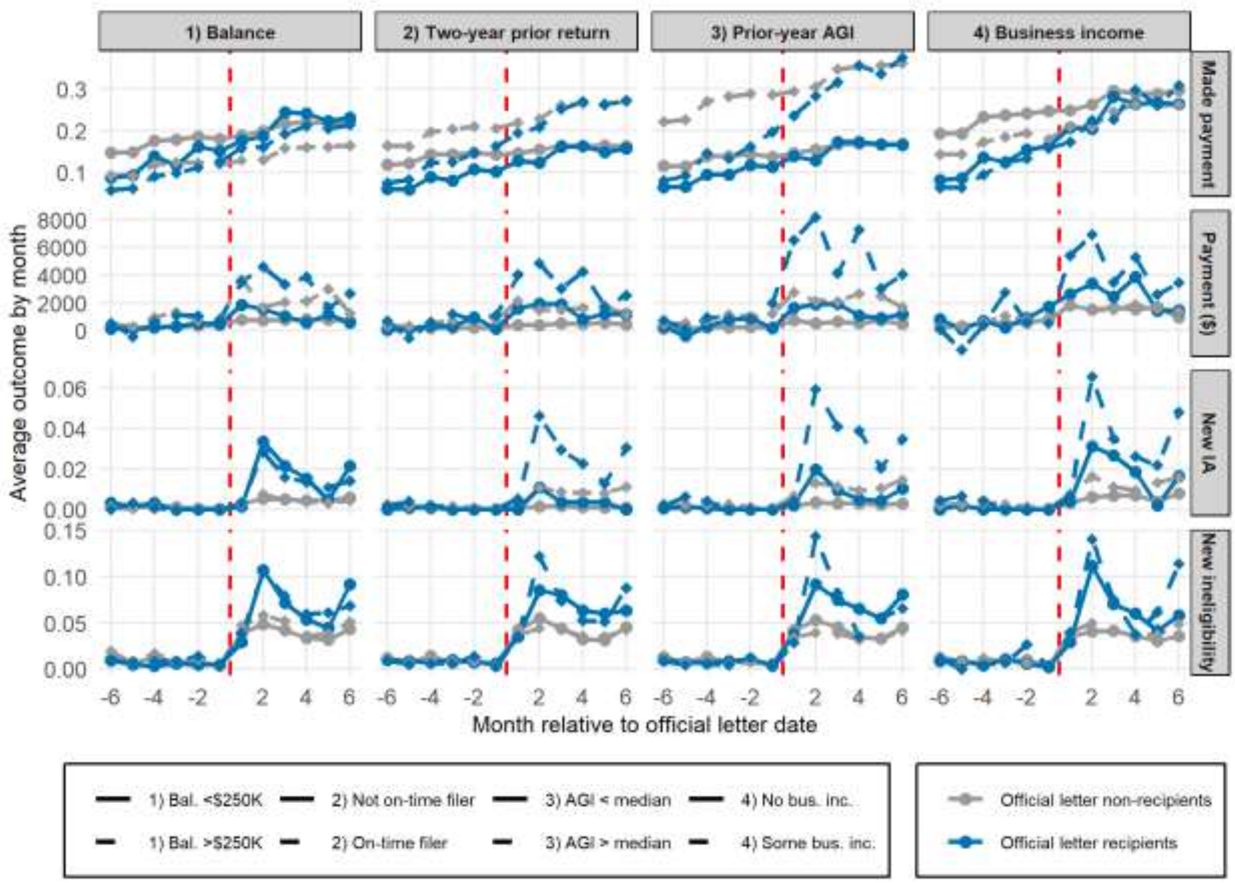

Notes: This figure compares the behavior of a two x four matrix of taxpayers around the date of the official letter. In gray are taxpayers who are eligible for publication and received a pre-letter, but did not receive an official letter. In blue are similar individuals (eligible pre-letter recipients) who did receive an official letter. In this figure we exclude individuals who previously received an official letter (i.e., we focus on first-time recipients). For each of four outcomes on the Y axis there are four groups on the X axis, each divided into two sub-groups. Dashed lines represent one of the sub-groups, solid lines the other.

We next examine these relationships in a regression framework. We repeat the analysis both restricting to households with balances within the cutoff range, and also over the full sample. The estimates thus take the form:

$$Outcome_{itc} = \alpha + \beta_1 \cdot Post_{itc} + \beta_2 \cdot Treat_{itc} + \beta_3 \cdot Post * Treat_{itc} + \beta_4 \cdot Characteristic_{itc} + \beta_5 \cdot Post * Characteristic_{itc} + \beta_6 \cdot Post * Treat_{itc} * Characteristic_{itc} + \gamma \cdot Balance_{ic} + April_c + \varepsilon_{itc}$$

where *Characteristic* is one of the four sub-groupings by balance, AGI, filing status, and business income. The coefficient of interest is β_6, the continuous incremental effect of treatment per unit of AGI (in millions here, for coefficient comparability), or the discrete incremental effect of having above-median AGI, non-negative AGI, the presence of business income, or on-time filing. Regression results are summarized in Figure 3.7 and tabulated in more detail in the Appendix.

Figure 3.7: Estimates of treatment by sub-groups

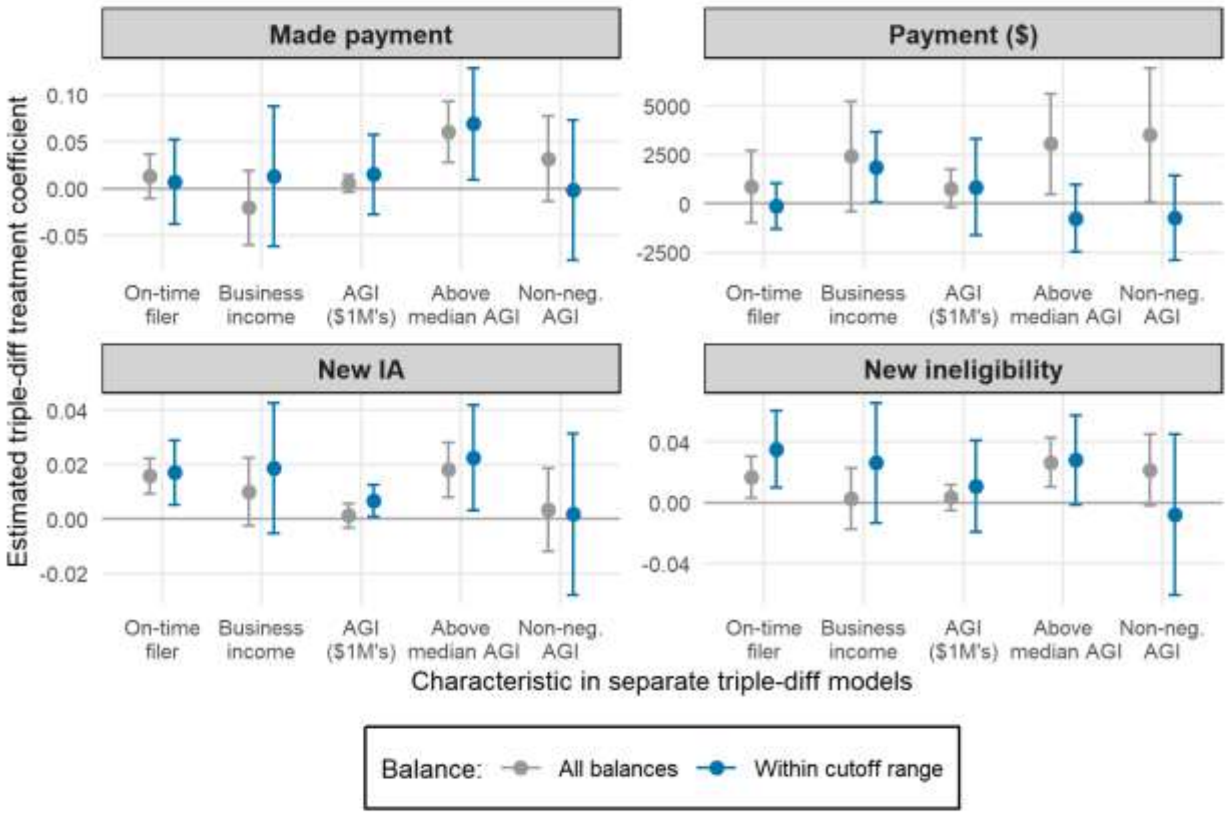

Notes: This figure presents the coefficient estimates for separate triple difference models testing four characteristics (i.e., the estimate for the coefficient on treatment X post X characteristic). 95% confidence intervals are shown around the point estimates. In blue are the estimates using only observations within the cutoff range. These can be compared to the estimates in gray, from models using the full range of balance observations. Corresponding tables can be found in the Appendix.

Income level appears to play an important role in responses to treatment. Treated households with above-median AGI are more likely to make a payment, enter into an Installment Agreement, or otherwise establish ineligibility.[87] These effects become insignificant when considering non-negative AGI instead of above median AGI, suggesting it is indeed higher AGI amounts that are driving this effect. In the full sample, average monthly payments after treatment are also much

[87] We find no significant effects of an interaction between treatment and a linear and continuous measure of AGI. This is not surprising, as there is no particular reason to expect that the impact of AGI will be linear in AGI. As an alternative, we also include interactions with AGI quintiles. Although less precisely measured, results for the upper quintiles, particularly the topmost quintile, are similar to those for the above-median results we present in the main text (see Appendix Figure 3.19).

higher among households with above-median AGI, though our point estimate is close to zero when we limit only to balances in the historic cutoff range.

In addition, we see a relatively large and statistically significant increase in the impact of treatment for filers with business income, with a point estimate about ten times larger than for all filers, although still moderately sized in economic terms, about $1,900 in additional payments each month. To be sure, we can only observe these outcomes for the subset of households for which we have tax filings, and so they may not be fully representative of all households. What we can say for certain is that conditional on filing, business income predicts a greater response.

We do not find evidence that filing status affects payment, but we do find it correlated with other compliance outcomes. For average monthly payment, our point estimate for the interaction of filing status with treatment is small and negative. Our confidence interval is fairly wide, however, so that we cannot rule out increases or decreases of $1,000. On-time filing does strongly correlate with an increased propensity to reach an Installment Agreement or other ineligible status, with 95% confidence intervals suggesting that these are both about twice as common among the filer population as among other treated households.

3.4.5. Long-run results

For purposes of program evaluation, it is useful to know whether treatment leads official letter recipients to remit more money than others over the long term. Among other reasons, one of the main responses we observe is a greater share of taxpayers who enter into payment agreements with FTB. Do these agreements actually bring in more money over time? At a minimum, it would be useful to know whether individuals who reach agreements in order to avoid publication quickly renege.

We can readily rule out the possibility that installment agreements are quickly broken. We sum payments by official letter-recipient households, and compare those who signed new installment agreements within three months of receiving a letter against all other recipients, as summarized in Figure 3.8. Installment Agreements strongly predict increased payments, whether over six, twelve, or twenty-four months after treatment. This result holds among all recipients, and also when restricting to balances within the cutoff range

Figure 3.8: Payment share among official letter recipients, effect of IAs

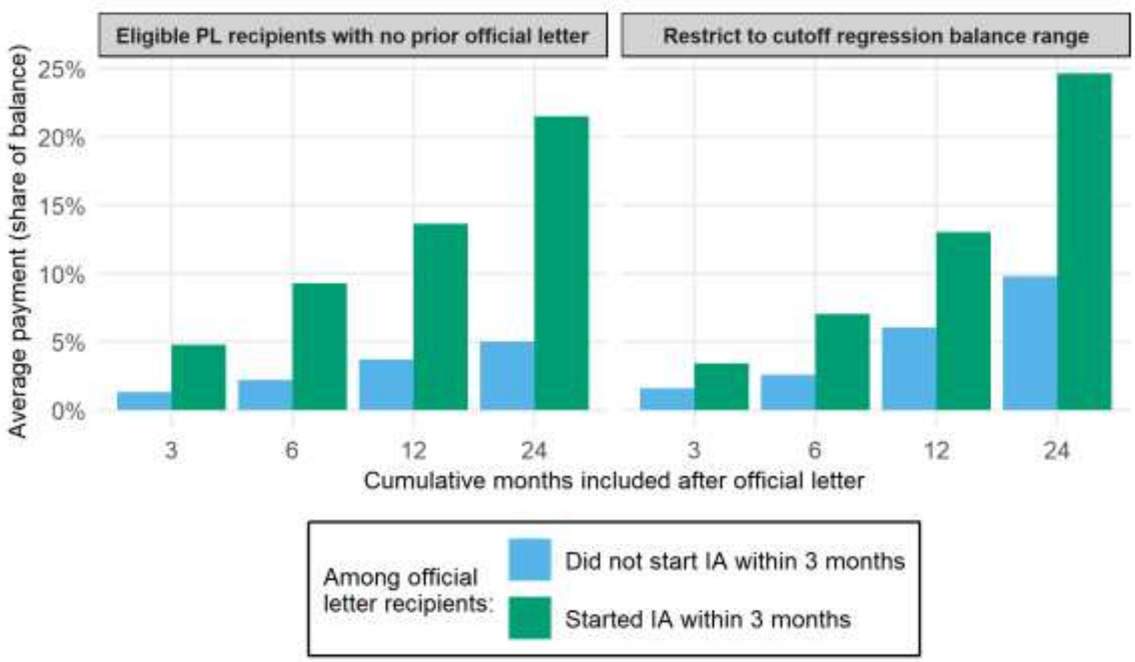

Notes: The figure summarizes average payments as a share of starting balance by households entering an installment agreement with FTB within three months of receiving an official letter, in green. Blue bars represent averages for all other treated households. The sub-graph on the right restricts observations to those where the unpaid balance at time of the letter falls within the cutoff range.

To provide at least a suggestive sense of other longer-run impacts of treatment, we repeat our regression analyses above, but collect cumulative results over the one- and two-year periods following a letter cycle. Admittedly, though, as we extend our observation window over a longer horizon, it is harder to rule out confounding effects. We thus present just a figure summarizing the main coefficients of interest for most of the outcomes.

Figure 3.9: Official letter long-term effect estimates

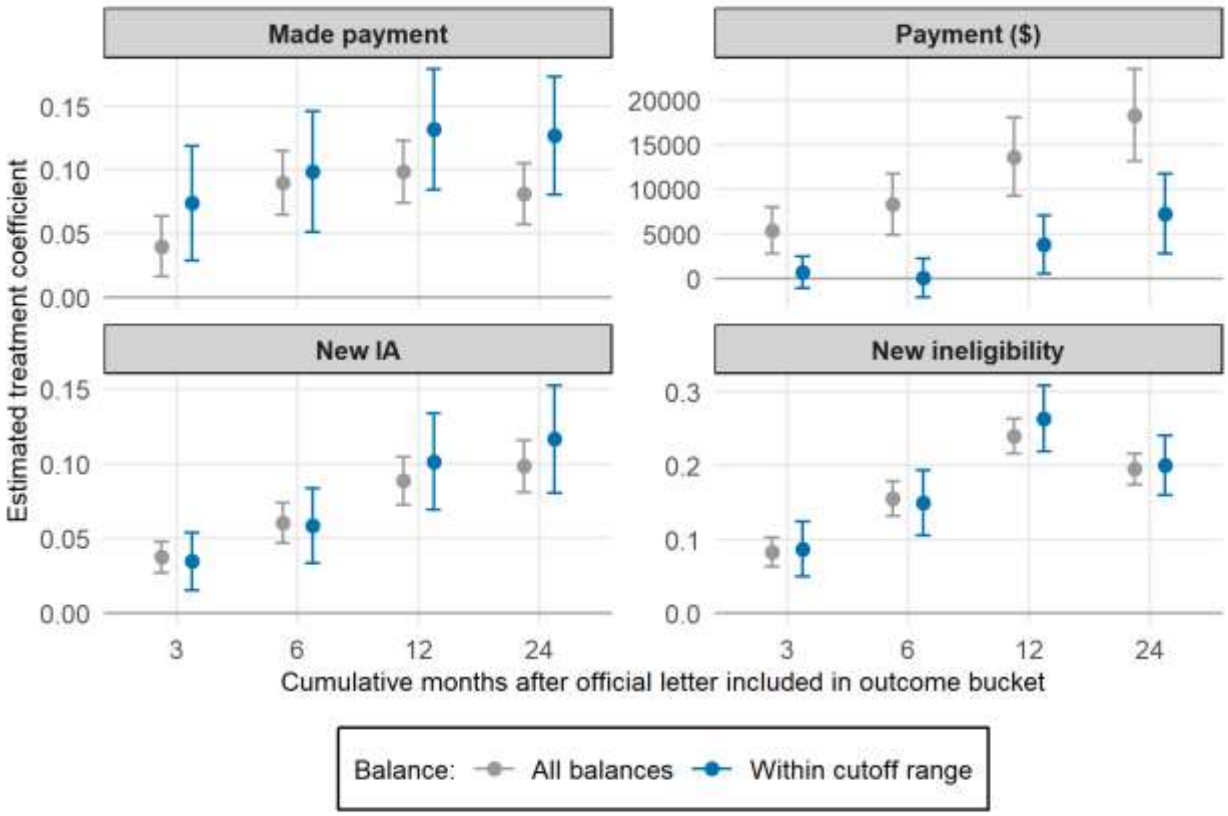

Notes: This figure summarizes estimated coefficients for regressions in which the variable in the grey bar is the outcome, cumulatively defined over the number of months post-official letter noted on the X-axis. Blue markers are for regressions in which we restrict the sample to observations with balances falling in between the lowest and highest observed Top 500 cutoff balances.

Because total payments are of particular interest for our evaluation framework, we present more detail on the long-run impact of treatment on total revenue collected. Although, again, the short-run effects of treatment are not statistically significant when we restrict to between-cutoff balances, over time these units do pay a good bit more, with a point estimate for incremental payments over two years of about $7,200 per household, and a 95% confidence interval ranging from $2,800 to $11,700, as shown in Table 3.4 below. As above, when we include all households in the analysis, the estimate is again much larger, with a point estimate of about $18,000 and a 95% confidence interval of $12,800 to $23,100 (see Appendix Table 3.8). As one further alternative approach, we run these tests excluding those we deem "partial controls" – individuals in the control group who receive a letter in a subsequent cycle and whose behavior during the outcome window may thus reflect a response to that later letter.[88] The estimated effects under this alternative approach are larger, although less precisely estimated (Appendix Figure 3.20).

Table 3.4: Official letter long-term payment effects (observations within cutoff range)

	Dependent variables: Cumulative payment amount post-official-letter ($)			
	3 months	6 months	12 months	24 months
Official letter	**693.15**	**111.26**	**3833.59****	**7260.01*****
	(904.73)	**(1,099.56)**	**(1,658.15)**	**(2,274.17)**
Balance	-0.19	5.89	-10.52	-27.72
($ thousands)	(9.65)	(14.86)	(20.24)	(29.34)
April publication	-1104.84***	-839.11	-25.76	335.50
	(414.60)	(612.57)	(680.92)	(824.58)
Intercept	3064.36*	4472.92*	10583.07***	18379.12***
	(1,700.09)	(2,708.97)	(3,710.08)	(5,332.18)
Observations	6,308	6,308	6,308	6,308
R2	0.0013	0.0003	0.0010	0.0021
Mean dep var.	2,539	5,126	8,980	14,129

Notes: *p<0.1; **p<0.05; ***p<0.01. Standard errors clustered by taxpayer are shown in parentheses. Outcomes measured as cumulative payments between time of treatment and time following, as listed in column headers. Only observations within the cutoff range are included.

3.4.6. Subsequent reported earnings

As we noted earlier, it is possible that disclosure programs are counter-productive if they reduce taxpayer ability to pay in the long run, as findings by DT suggested might be the case for some businesses. Likewise, disclosure might backfire if it crowds out future voluntary compliance efforts, potentially reducing *reported* income. Accordingly, we exploit our ability to link payments

[88] Although retaining some partially-treated units could bias our results downwards, we prefer those estimates because we view them as more conservative. Dropping taxpayers who are subsequently treated from the control could potentially bias results upwards. Taxpayers who are treated in a subsequent cycle are by definition non-compliant for an extended period, and may even have taken actions that deepen their debt. Comparing our treated group to this selected non-responsive group might overstate the effects of treatment relative to the general population of tax delinquents.

data to tax filings to test whether there are any observable long-run impacts on income reported to FTB.

In general, although our point estimates are positive and relatively large, we cannot rule out economically meaningful declines in reported earnings. In our full sample, for instance, the 95% confidence interval for the household's change in AGI between treatment and two years after treatment ranges from -$128,000 to $435,000. Because we have access only to pre-audited income, we cannot tell whether any possible declines might be due to actual reductions in taxpayer earnings or whether they are only a change in reporting.

Figure 3.10: Effects of official letter on subsequent AGI

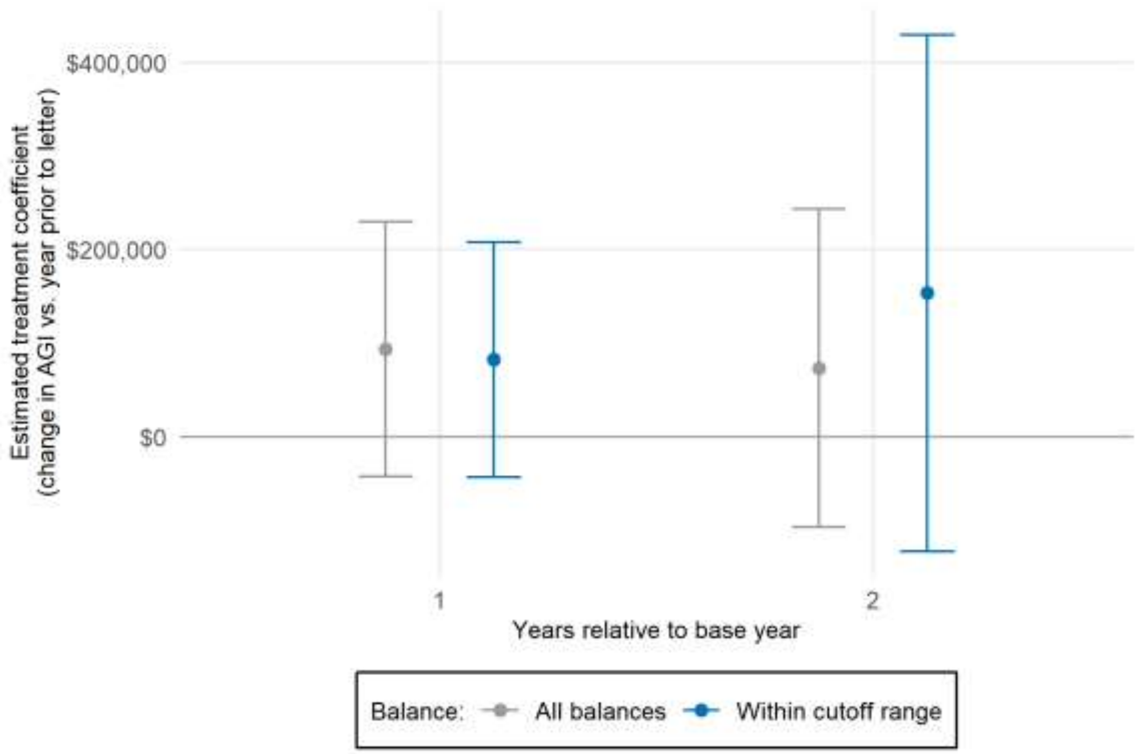

Notes: This figure presents the coefficient estimates and 95% confidence intervals for regressions testing the effect of official letter receipt on subsequent changes in AGI.

3.5. Publication and license revocation

The previous section showed that notifying taxpayers that their information will be published if they do not take action clearly causes a substantial number to take action to avoid publication. Next, we attempt to learn whether publication itself has any effect on the taxpayers who ultimately do get published. As with the official letter, there is potential selection before publication, as the individuals who are published on the Top 500 list are the individuals who were given an opportunity to avoid publication, by taking action after the official letter, *and chose not to take such action*. As we do observe a fair bit of response to the official letter, we may not expect publication itself to have much additional impact. DT do find a moderate incremental effect from publication, but their study involved a first-time rollout of a program, such that it may not have been clear *ex ante* whether the government would carry through with its threat or what impact disclosure would have.

Because we aim at the marginal compliance effect of publication over and above receipt of the official letter, we restrict our analysis here to official letter recipients. We then observe the behavior of the two groups before and after publication. As above, we focus on first-time letter recipients to understand the effect of a taxpayer's first encounter with the risk of publication.[89]

Figure 3.11 below shows the time series comparison of the published and unpublished first-time official letter recipients. The red vertical line indicates the official letter date, while the blue vertical line indicates the Top 500 publication date. The gray series represents the average behavior among the first-time official letter recipients who are *not* published. As expected, we observe larger spikes at the time of the official letter for this group than we plotted in earlier figures, as we are splitting out the subgroup that did not respond.

[89] The findings in this section are similar if we include all official letter recipients, as shown in the Appendix.

Figure 3.11: Behavior around the publication date among first-time official letter recipients

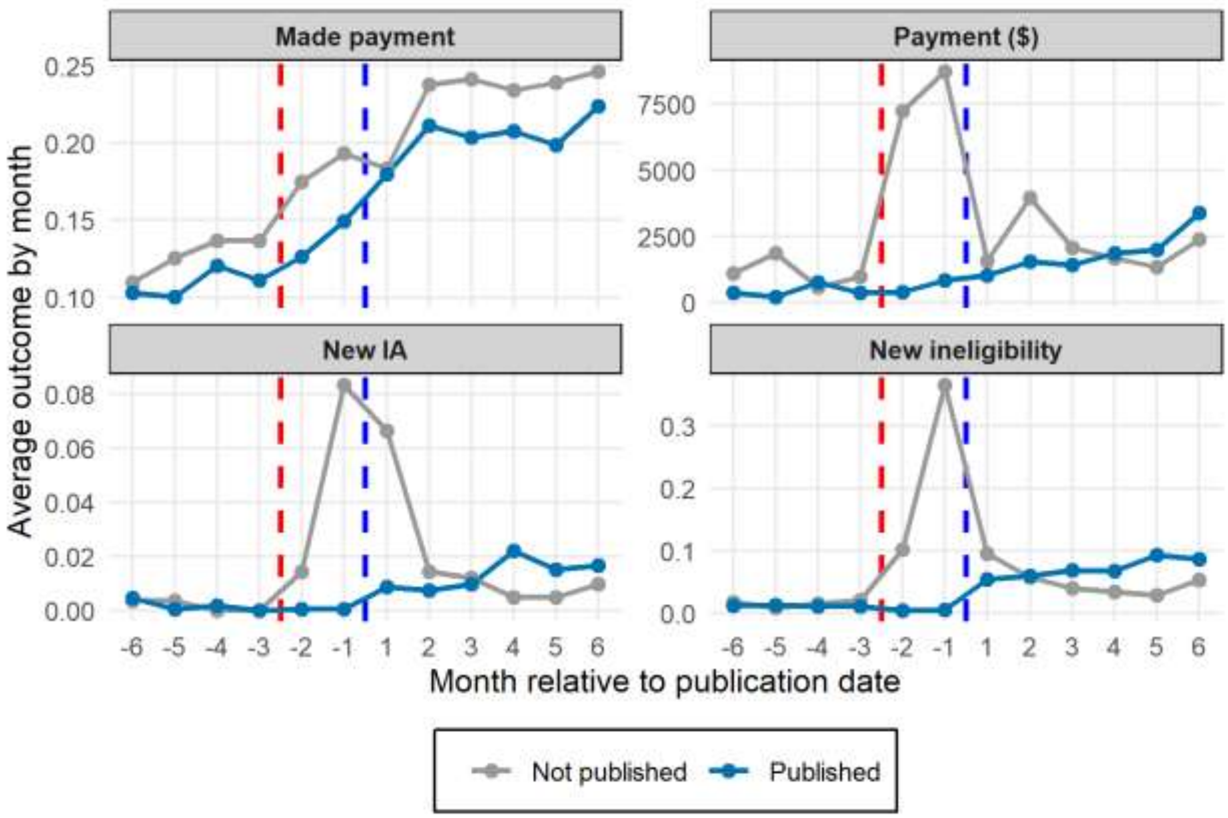

Notes: This figure compares the behavior of first-time official letter recipients, separately showing those that ultimately get published (in light blue) and do not (in gray). The red vertical line indicates the official letter date, and the blue vertical line indicates the Top 500 publication date.

The light blue series represents the average behavior among the first-time official letter recipients who *do* get published. Their lack of response to the letter is what leads them to be published. We see little evidence that this group responds to publication, except for a small bump in new installment agreements about four months after publication. This timing coincides with the statutory timing for license revocation notices, which issue 120 days after publication (after an initial warning 30 days following).

To better understand whether the observed response may be related to license suspension, we further segregate the population by whether FTB data show any professional license that might be subject to suspension. FTB only collects these data for published taxpayers, so we cannot provide a full triple-differences analysis.[90] Graphical analysis is suggestive, though, that license suspension does have some impact, as illustrated in Figure 3.12.

Figure 3.12: Behavior around the publication date, split by license holding

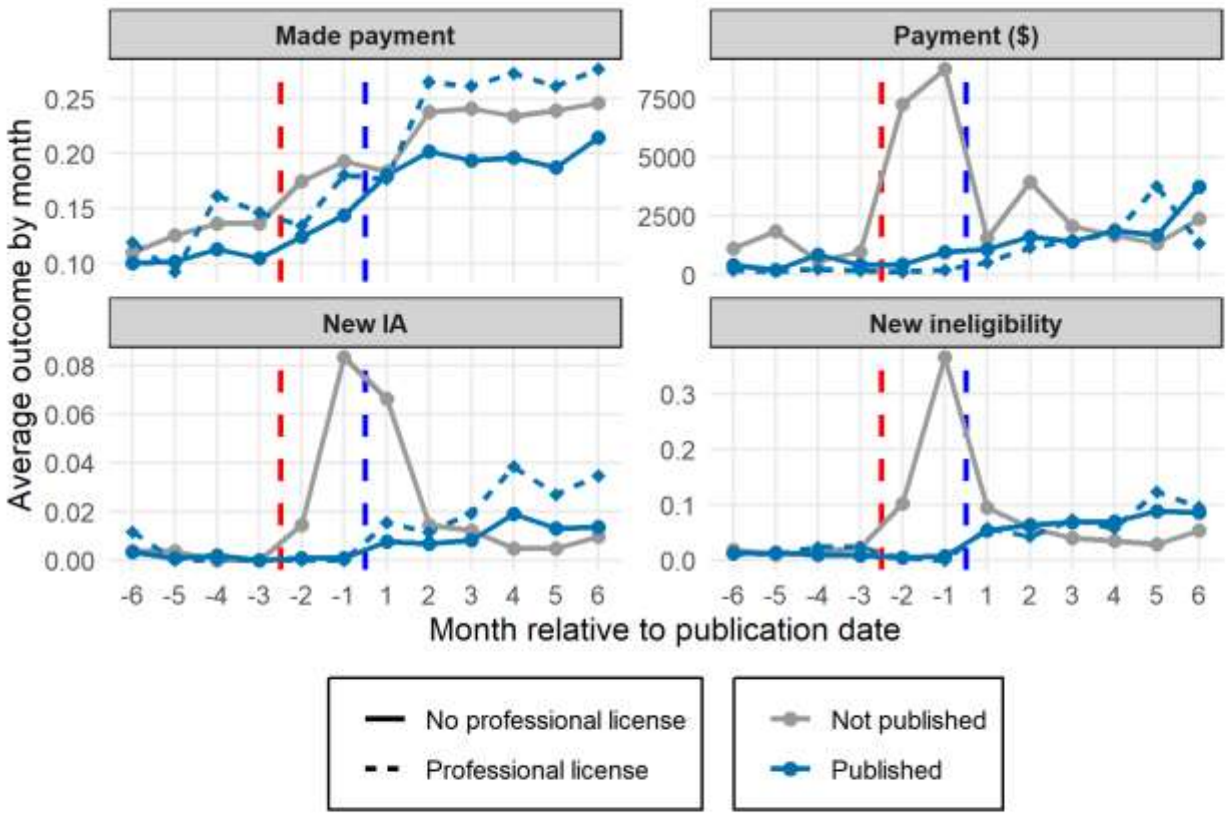

Notes: This figure compares the behavior of first-time official letter recipients, separately showing those that ultimately get published (in light blue) and do not (in gray). The published individuals are further split into those with professional licenses (dashed line) and those without (solid line). The red vertical line indicates the official letter date, and the blue vertical line indicates the Top 500 publication date.

In the figure, there is a noticeable above-trend surge in the likelihood of making any payment for license holders (plotted using the dashed light blue line) around the date of the first license notification, 30 days after publication. We also see a slightly higher share of installment agreements, peaking at the time when license suspension would take effect, 120 days after publication.

[90] If data on licenses were available for all individuals, it would be interesting to test for differential responses to the official letter based on license-holding and among different license types. It is possible that some of the response to the official letter is driven by concern over future license suspension, rather than publication.

3.6. An evaluation framework for tax debt collection and non-monetary sanctions

With these outcomes in hand, we now aim to evaluate the California program through two related frameworks. First, we ask whether taken on its own it likely increases social welfare, relative to a baseline of no added enforcement of any kind. Second, we ask whether disclosure is optimal given alternative supplemental enforcement choices, such as increased fines or penalties on late payers.

Analysis of the first frame is familiar. Tax compliance efforts are not universally welfare-improving. Instead, as Keen and Slemrod (2017) show, the necessary condition for welfare-improving compliance policy α is:

$$\phi(tz_\alpha - a_\alpha) - c_\alpha > 0 \tag{1}$$

where tz_α is total tax revenue caused by the policy (tax rate t times marginal taxable income z_α), a_α is the administrative cost of the policy, and c_α is the net marginal compliance or concealment cost. ϕ is the weight applied to government revenues, generally the marginal social value of public spending (Meiselman 2018), which we assume to equal the marginal cost of public funds (Hendren and Sprung-Keyser 2020 provides a more comprehensive estimate of MVPEs for an array of spending options). In words, a public expenditure on increased compliance increases welfare when the marginal value of additional public funds, net of marginal enforcement, avoidance, and compliance costs, exceeds zero.

Net private costs c_α can be either positive or negative. Intuitively, when some tax avoiders become compliant, they no longer incur private avoidance costs, but instead must bear the costs of compliance, while infra-marginal avoiders may strictly increase avoidance expenditures. In a simple setting where taxes are not shifted across bases or time, we can infer that, for taxpayers at the margin, the amount of additional tax paid and compliance costs incurred equals the amount of private avoidance costs saved (Feldstein 1999, Chetty 2009). Thus, if marginal compliance costs are small, the marginal revenues from an enforcement effort, weighted by the value of public expenditures, and net of the public expenditures on that effort, offer a reasonable starting point for the social benefit of enforcement (Keen and Slemrod 2017). In our setting, direct compliance costs over and above those of payment itself are likely minimal.[91] But marginal private costs should also include any measurable impact on avoidance expenditures by infra-marginal avoiders.

Accordingly, we aim to measure what we argue are the two key inputs into this basic welfare analysis: net revenues and infra-marginal avoidance costs. Social planners could then weight our net revenue estimate by their preferred value for the marginal value of public expenditures to assess whether the Top 500 program increases welfare relative to a baseline of no special enforcement policies for the largest debtors.

It is also useful to evaluate whether use of the Top 500 program is preferable to alternative methods for collecting large tax debts. For example, standard enforcement theory holds that monetary sanctions are usually preferable to non-monetary sanctions such as disclosure (Polinsky and Shavell 2000). While both might be capable of motivating compliance, a non-monetary sanction such as disclosure imposes costs on non-compliers with no offsetting gains, whereas a fine is a transfer and potentially welfare-neutral.[92]

More recent work suggests at least three potential reasons for preferring disclosure in select instances. The first of these relates to the possible effect of disclosure on tax evasion, as in

[91] To the extent that the increased threat of collection discourages avoidance or evasion efforts prior to assessment, collection may affect income shifting and compliance expenses. Our discussion in this paragraph focuses on measurement of the direct social gains from collections, not these kinds of "upstream" effects.

[92] The net welfare effects of a transfer may depend on the relative social welfare weights assigned to transferor and transferee. For example, Galle (2013) discusses informally the importance of the relative weights on transferors and transferees in comparing alternative Pigouvian instruments. Kaplow (2008) argues, however, that transfers should always be considered to be welfare-neutral because any unwanted redistribution can be undone by, or desirable distribution accomplished through, income tax schedules.

Kuchumova (2018), which models both license and passport suspension as instruments for reducing evasion. In her framework, the non-monetary sanction potentially serves as a "tag" for higher true earnings. Because driving and international travel are forms of consumption that are correlated with income but largely unobservable by the tax authority, suspension effectively imposes a higher tax on individuals with greater true earnings. Welfare gains from this improved targeting can exceed welfare losses from the deadweight loss the transferless instrument imposes, as in equation (2):

$$W_1 - W_2 > c_n - c_m \tag{2}$$

which corresponds to Kuchumova's equation (26), but where we have simplified to set W_1 - W_2 as the net social gain caused by transfers to low potential earners from high earners, and c_n - c_m as the incremental compliance costs of the nonmonetary instrument over the monetary instrument. This is a threshold that the non-monetary policy must clear in addition to equation (1). That is, if both non-monetary and monetary sanctions would satisfy cost-benefit analysis, the non-monetary sanction is only optimal if the incremental transfers achievable through the non-monetary instrument exceed its additional deadweight loss. We expect c_n - c_m to be hard to estimate because it will often require comparison to a hypothetical monetary instrument.

This same model can extend to disclosure if it is the case that the evader's taste for being perceived as tax compliant is also valued more highly by those with higher true earnings. DT report, for example, that larger firms in their sample were more responsive to treatment. We test this hypothesis with respect to individuals.

A second possibility is that public disclosure is an efficient component of a "tax systems" approach to compliance (Slemrod and Gillitzer 2013). In their account, taxpayers have a menu of options for minimizing the household tax burden. Government responses may affect each of the taxpayers' margins differently, resulting in varying elasticities of taxpayer response to each government strategy. This results in a Ramsey-type model in which government should employ a variety of enforcement techniques, each weighted inversely to the elasticity of taxpayer response.

Keen and Slemrod (2017) extend the tax systems approach, showing that a similar inverse-elasticity rule holds in the presence of transfers. Extrapolating from their parts 4.1 and 4.2, we can say that:

$$E(z,\alpha_k) = (\alpha_k((1\text{-}\mu_k)\ c_{\alpha_k}\ /\ \phi) + \alpha_k a_{\alpha_k})\ /\ tz \qquad (3)$$

where $E(z,\alpha_k)$ is the elasticity of revenue with respect to the enforcement instrument α_k, $\alpha_k a_{\alpha_k}$ is the cost of administering that instrument, and $\alpha_k((1\text{-}\mu_k)\ c_{\alpha_k}\ /\ \phi)$ is the net cost of compliance, discounted by the proportion $1\text{-}\mu_k$ (with $0 \leq \mu_k \leq 1$) that represents the share of costs that are other than pure transfers. As above, ϕ is the marginal value of public spending and tz is revenue. That is, for any given instrument, the government should invest in enforcement to the point at which the marginal after-transfer cost-to-revenue ratio is equal to the enforcement elasticity of tax revenue.

Equation (3) suggests that non-monetary instruments can be optimal, but usually only if they are highly effective in returning revenue. A non-monetary instrument is likely to be optimal only when the elasticity of revenue with respect to enforcement is relatively high, as the top term on the right-hand side will be relatively large. In order for this to be the case, the non-monetary sanction would presumably have to affect different margins of response than the monetary sanction, or have a larger elasticity of response per unit of expenditure; otherwise, increasing the monetary sanction would strictly dominate (again, ignoring welfare weights). As Galle and Mungan (2021) show in a Pigouvian setting, it is usually optimal to exhaust monetary sanctions before employing non-monetary options, but non-monetary options can still be optimal when taxpayers are heterogeneous in their sensitivity to sanctions.

Taken together, then, this second set of prior studies suggests that disclosure could potentially offer an efficient tool for tax collection. Households with high ability to shield wealth from collections are unlikely to respond to a threat of fines that would be uncollectable. If these taxpayers were not sanctioned, they would have an *ex ante* incentive to shift to uncollectable sources of income. But households may find it more difficult to escape disclosure than a fine, reducing their propensity to earn uncollectable income.

Kuchumova sets out a similar model, but focused on the social welfare benefits from collecting assessed taxes rather than the deadweight loss from taxpayer's "upstream" decisions about sources of income (Kuchumova 2021). As in earlier work (Andreoni 1992), she observes that a taxpayer's ability to pay a debt at the time of collection gives us more information about their marginal utility; taxpayers with low ability to pay should not be subject to collection actions because these households have very high marginal utility of a dollar.[93] The difficulty is that if households can also hide their wealth, low-marginal-utility households can mimic high. Non-monetary sanctions fall more heavily on mimic households than genuine high-utility taxpayers, however, and so a separating equilibrium becomes possible: the mimic household can delay a fine by hiding household wealth, but cannot similarly delay non-monetary sanctions. Kuchumova shows that the resulting gains from this strategy can outweigh the deadweight loss of relying on transferless instruments as long as the population of debtors includes a sufficiently large share of

[93] We accept this point for our limited purposes here. We observe, however, that forbearing on tax debts is unlikely to be the government's best instrument for redistributing to low-income households, as it has effects on tax morale and upstream tax planning that alternatives likely do not.

high-income households, and this threshold becomes easier to meet as the social value of tax revenue (or, equivalently, the social cost of collecting substitute revenues) rises.

In addition to these three direct effects on tax debtors, disclosure systems in particular may affect taxpayer morale. By highlighting state efforts to ensure that everyone pays their fair share, disclosure may contribute to a sense that tax systems are fair, and thereby encourage voluntary compliance (Blank 2014; see Tyler 1999 for evidence that perceptions of systemic fairness affect compliance generally, and Hartner, et al. 2011 and Murphy 2003 for evidence in the tax context). Disclosure may also backfire, however, if it instead spotlights wrongdoing or otherwise causes observers to update their beliefs about compliance downwards (Luttmer and Singhal 2014).

In short, although non-monetary sanctions pose tradeoffs, they also offer a path to a more efficient tax system. Households with the ability to conceal collectable wealth may be relatively insensitive to additional fines or fees, but still relatively sensitive to disclosure or a lost license. These instruments may therefore both bring in assessed revenues as well as deter behaviors that would prevent the revenues from being assessed in the first instance.

3.7. Application of the evaluation framework

To apply the frameworks just set out, we use our estimates in Sections 4 and 5 to pin down values for equation (1), our measure of the social welfare effects of the program, relative to a baseline in which there is no alternative enforcement policy. For easy reference, we repeat equation (1) here in words:

$$\phi(Revenues - Admin\ Costs) - Private\ Compliance\ Costs + Foregone\ Avoidance\ Costs > 0$$

Our estimates suggest the California Top 500 program brings in meaningful amounts of revenue. Our preferred point estimate for the incremental two-year payments of first-time treated households is about $7,200, and there are an average of 400 such households per year. A simple back-of-the-envelope calculation thus suggests the program results directly in at least $2.8 million annually in its steady state. We argue, though, that this figure is too low, because it fails to account for higher-balance households, where most of the outstanding debt is. Although our estimates for these households are not as well-controlled, it is likely that any selection that occurs actually depresses our estimate. When we use estimates for all eligible households with a balance above $100,000, the back-of-the-envelope revenue figure is $7.2 million.[94]

An additional benefit from inducing taxpayer compliance is that taxpayers do not incur deadweight-loss costs of avoiding collection. As discussed, we argue that in a rational-actor sufficient statistics framework, increased payments of $7.2 million in this setting imply as much as $7.2 million in foregone private avoidance.

On the cost side, it is not easy to fully separate the costs of the program from general collection costs, but direct estimated administrative costs total somewhere between $1.5 and $2 million per year. This reflects the estimated direct cost of administering both the personal income tax list and the corporate income tax list; the personal income tax cases typically represent about 80% of the total, so the relevant cost for this study is in the range of $1.2 to $1.6 million. To ensure a conservative evaluation of the program, we will use the $1.6 million estimate.

Estimating private compliance costs is a bit more subtle. Non-compliers who are unwilling or unable to pay are posted to the list, and our evidence suggests that at least the complier portion of the population perceives publication to be subjectively costly. Holding all else equal, the revealed preference of non-compliers with ability to pay is that their disutility from publication is lower

[94] We emphasize that this figure represents only the estimated mean incremental payments caused by treatment among individuals receiving the official letter in the two years after receipt, and so ignores payments by the thousands of other recipients of the pre-letter, as well as omitting payments by business entities. Total annual revenues brought in by the Top 500 unit are at least an order of magnitude larger, although the exact figure is not public information.

than that of compliers. Subjective costs among non-compliers are thus likely lower on average than among compliers.

We therefore can put a plausible upper bound on the disutility from publication experienced by non-compliers. We infer that the disutility per household for non-compliers is strictly lower than for compliers. Among compliers, the observable cost of avoiding publication is the amount of incremental payments these households make over our two-year window, approximately $46,000 in the full-sample results. Over the ten observed cycles, we count an average of 161 initial non-compliers per cycle, defined as households that are published for the first time. This yields a ceiling value of $7.4 million.

Some non-compliers may fail to comply because they lack the ability to pay, but this doesn't affect our cost calculation. Although we lack direct evidence of the subjective cost of disclosure for those who cannot pay, there is no reason to think it differs from the distribution in the general population (Goldin and Reck 2020). Certainly, there is no reason to think it would be higher on average than the average disutility among compliers—higher, that is, than among households that reveal themselves to have the greatest subjective costs of publication. Thus, if we assign a value of $46,000 per household to all households, we still have something of a ceiling on private costs: $46,000 is the mean revealed value among compliers, is the maximum for intentional non-compliers, and is likely the mean value for those who lack ability to pay.

To reiterate, when drawing on our full sample results we estimate additional revenue of $7.2m, administrative costs of $1.6m, private compliance costs of $7.4m in disutility of publication and $7.2m in payments (a total of $14.6m), and foregone avoidance costs of $7.2m.[95] This results in net social welfare from the program of approximately:

[95] Arguably, our calculation of private compliance costs should omit the cost of payments, but we include them to present a conservative estimate. Typically, measures of the efficiency of the tax system will already account for the

$$\phi(\$7.2m - \$1.6m) - \$14.6m + \$7.2m = (\phi)\$5.6m - \$7.4m$$

If we assume that the marginal value of public revenue is equivalent to the tax cost of replacement, and use common estimates of the (national) MCPF of 1.5 (Heckman et al. 2010; Cellini et al. 2010 also estimate the marginal value of public expenditures in California at 1.5), we estimate total annual welfare gains of about $1 million, ignoring any possible distributive weights. Again, this estimate relies on an assumption that the actors we observe are rational and on the margin between compliance and non-compliance. The ultimate value could be higher (lower) if a larger (smaller) share of non-compliers are infra-marginal as compared to compliers.

We note that netting incremental revenue collection against these private costs likely considerably understates the benefit of the Top 500 program, because we cannot directly measure its upstream impacts on taxpayer behavior. For instance, we do not observe the extent to which the program may contribute to taxpayer morale.

Our results also shed some light on the second possible evaluation framework, namely, whether the Top 500 program is more efficient than other alternative interventions. Testing equation (2), derived from Kuchumova 2018, requires us to make assumptions about the efficacy of a counter-factual regime that relied solely on elevated fines and fees. Given that the treated households we observe had already avoided paying hundreds of thousands of dollars in debts, we think it is reasonable to assume additional fines and fees would have had limited effect. If so, equation (2) is likely satisfied. We can take $\phi(Revenues - Admin\ Costs)$ as the weighted social value of added revenue derived from tagging. And the deadweight loss of disclosure is a fair

private costs of payment (Dahlby 2008). If we are estimating the incremental welfare effects of a particular tax collection method, accounting for the costs of payment would thus be double-counting. If payments are omitted from the private compliance costs, our estimate of net social welfare increases to $8.2m.

estimate for c_n - c_m, the incremental compliance costs of the Top 500 program. This again gives us an estimate in the range of $1m.

We also can make some observations about equation (3), which requires that the weighted cost-revenue ratio exceed the elasticity of income with respect to the enforcement tool. Plugging our point estimates into the right-hand side of the equation yields:

((($7.4m - $7.2m) / 1.5) + $1.6m) / $7.2m ≈ 0.24

This calculation reflects the net of compliance costs ($7.4m in deadweight loss, none of which is transferred to others, less $7.2m in foregone avoidance expenses), plus the conservatively high estimate of $1.6m in administrative costs, over our revenue point estimate of $7.2m.

We cannot fully evaluate this result without knowing the opportunity set available to FTB. In general, though, a cost-revenue ratio this small implies that the Top 500 program would be an efficient option even if its elasticity of revenue with respect to enforcement were very low.

Additionally, non-monetary sanctions produce welfare gains if they affect a population that would be able to escape taxation if it were enforced only through a fine. We find a large relative increase in payments by households with business income, suggesting that the implementation of a disclosure regime helps to reduce the net tax-avoidance payoff to self-employment. To be sure, this mechanism is imperfect: a fairly large fraction of delinquent households did not file any recent return, and non-filers are relatively less likely to respond to treatment. While this could reflect relatively high subjective costs of compliance, it may also reflect evasion behavior that is not curbed by disclosure.

Our results shed light not only on whether the Top 500 Program passes cost-benefit analysis, but also more generally on why households appear on the Top 500 list. As expected, ability to pay is part of the story. We see that, holding balance constant, households with the highest reported

AGI respond to treatment with the greatest payments, the largest increase in likelihood of making any payment, and the greatest propensity to enter into an installment agreement.

But ability to pay is not the full story. If households are able to respond to treatment by paying tens of thousands of dollars on average, why don't they pay when their taxes are due, or when they receive personal contact from an FTB employee, or when they receive a notice of tax lien? At prevailing California income tax rates, a taxpayer would likely have needed to earn at some point in the past in aggregate hundreds of thousands of dollars in taxable (not gross) income in order to accumulate tax debts in excess of our lowest observed list cutoff value of $150,000. Income volatility may result in large fluctuations in ability to pay over time, but would not explain why treated households are more apt to pay than others. We further see a strong correlation of basic tax compliance through the filing of a return with subsequent willingness to comply.

The data are therefore consistent with a substantial amount of deliberate tax avoidance among the observed population. As an additional piece of evidence on this front, we note that tax-filing business owners, who are not subject to withholding, are far more responsive to treatment than the average tax-filing household. This is consistent with the hypothesis that business owners are better able to protect their assets from other collection efforts. It is also possible that business-owners are more responsive because disclosure is more costly, such as by affecting the business's relationship with customers or suppliers (as noted by DT). We don't observe any significant differences in post-treatment reported earnings for business owners versus other taxpayers, however.

3.8. Conclusion

We have examined the economic impact of state disclosure of individual taxpayers with large unpaid tax debts. Households receiving a warning that disclosure is imminent respond with large increases, relative to baseline, in their efforts to reach payment agreements with the California tax authority. Households also pay more, with increases in payment particularly concentrated among those with greatest ability to pay and those with business income. Cost-benefit analysis suggests that, looking only at the most direct effects of the program, it increases social welfare substantially. We also find evidence that a good portion of this gain derives from reduced tax avoidance or evasion.

The California experience therefore suggests that delinquent taxpayer disclosure programs may be a useful component of a regulatory toolkit, particularly when compared to other non-monetary sanctions. The households we have studied are ones where the government has all but exhausted its ordinary collection efforts. That we identify a relatively sizable additional response from the threat of disclosure even from this population suggests that disclosure would likely be quite effective for more-compliant households, especially if such households also attach a higher subjective value to being *seen* as compliant. Our analysis also suggests some caution, though, in employing delinquency disclosure for households with lifetime incomes considerably lower than those we observe.

3.9. Appendix A: Sample FTB Documents

Figure 3.13: Sample pre-letter

Information Regarding Public Disclosure of Tax Delinquency

Notice Date:
Taxable Years:

Account Number:
Balance Due:
Pay By:

Revenue and Taxation Code (R&TC) Section 19195 directs the Franchise Tax Board (FTB) to publicly disclose a list of the 500 largest state income tax delinquencies. These delinquencies must total in excess of $100,000 and be subject to a recorded notice of state tax lien. We intend to post this list on our website at **ftb.ca.gov**.

Your account may qualify for this disclosure and Internet posting. If we determine that you are among the 500 largest tax delinquencies, we will send you a notice by certified mail advising you of the inclusion on this list of your name, address, and any occupational or professional licenses with status. At the time of publication, **your occupational, professional, and driver licenses issued by a California agency will be submitted for suspension**. You will be prohibited from contracting with any California state agency for the acquisition of goods or services. **Final determination of the top 500 names eligible for publication is pending confirmation of any resolutions or other qualifying circumstances for exclusion, including payment or arrangement for payment of tax liabilities.**

Pay your debt in full. You may be required to make payments electronically. Go to **ftb.ca.gov** and search for **mandatory epay**. If your estimated tax or extension payment exceeds $20,000 or your tax liability exceeds $80,000 for any taxable year beginning on or after January 1, 2009, you must make all future payments electronically, regardless of the taxable year. Payments made by other means will result in a penalty of 1 percent of the amount paid, unless your failure to pay was for reasonable cause and not willful neglect (R&TC Section 19011.5). If you are not required to pay electronically, enclose the above part of this notice and mail it with a check or money order for the total amount due payable to the Franchise Tax Board. Write your full name and account number on your payment. **Use the enclosed return envelope and mail to:** FRANCHISE TAX BOARD, PO BOX 3065, RANCHO CORDOVA, CA 95741-3065. No additional penalties or interest accrue on the existing liability if we receive full payment within 15 days of the notice date.

Call 888.426.8555 if you have questions, need assistance, think you do not owe this amount, paid the balance due, or filed bankruptcy.

Get FTB 1131, *Franchise Tax Board Privacy Notice,* at **ftb.ca.gov** and search for **privacy notice**.

Get FTB 1140, *Personal Income Tax Collections Information,* at **ftb.ca.gov** and search for **1140**.

Internet and Telephone Assistance
Website: **ftb.ca.gov**
Telephone: 888.426.8555 from within the United States
916.845.7874 from outside the United States
TTY/TDD: 800.822.6268 for persons with hearing or speech impairments

FTB 3703 PIT PC (REV 03-2012)

Figure 3.14: Sample official letter

Notice of Public Disclosure of Tax Delinquency

Notice Date:
Taxable Years:

Account Number:
Balance Due:
Pay By:

Revenue and Taxation Code (R&TC) Section 19195 directs the Franchise Tax Board (FTB) to publicly disclose a list of the 500 largest state income tax delinquencies. These delinquencies must total in excess of $100,000 and be subject to a recorded notice of state tax lien. We intend to post this list on our website at **ftb.ca.gov**.

Your account qualifies for this disclosure and Internet posting. If you do not pay your tax liability or take other action described below, we may add to a list posted on our website:

- **Your name and address.**
- Your occupational or professional licenses with type, status, and license numbers.
- The liened amount owed and the earliest date a notice of state tax lien was recorded.

Your inclusion on the list may lead to the denial or suspension of your licenses, including driver's licenses, under Business and Professions Code Section 494.5, and will preclude you from entering into contracts for the acquisition of goods or services with California state agencies under Contract Code Section 10295.4.

To avoid public disclosure of tax delinquency, you must do **one** of the following within 30 days of the notice date:

- **Pay your balance due. You may be required to make payments electronically.** Go to **ftb.ca.gov** and search for **mandatory epay**. If your estimated tax or extension payment exceeds $20,000 or your tax liability exceeds $80,000 for any taxable year beginning on or after January 1, 2009, **you must make all future payments electronically**, regardless of the taxable year. Payments made by other means will result in a penalty of 1 percent of the amount paid, unless your failure to pay was for reasonable cause and not willful neglect (R&TC Section 19011.5). If you are not required to pay electronically, enclose the above part of this notice and mail it with a check or money order for the total amount due payable to the Franchise Tax Board. Write your full name and account number on your payment. **Use the enclosed return envelope and mail to:** FRANCHISE TAX BOARD, PO BOX 3065, RANCHO CORDOVA, CA 95741-3065. No additional penalties or interest accrue on the existing liability if we receive full payment within 15 days of the notice date.
- **Arrange to pay your balance due**. To determine if you qualify for installment payments, call us at 888.426.8555.

Partial payment (even a reduction of the balance due below $100,000) will not preclude you from being on the list.
If your name appears on the list, FTB will continue to pursue collection actions. Call 888.426.8555 if you believe you should not be on the list, have questions, paid the balance due, made payment arrangements, otherwise resolved the balance due, think you do not owe the balance due, or filed bankruptcy.

Get FTB 1131, *Franchise Tax Board Privacy Notice,* at **ftb.ca.gov** and search for **privacy notice**.

Get FTB 1140, *Personal Income Tax Collections Information,* at **ftb.ca.gov** and search for **1140**.

Internet and Telephone Assistance
Website: **ftb.ca.gov**
Telephone: 888.426.8555 from within the United States
916.845.7874 from outside the United States
TTY/TDD: 800.822.6268 for persons with hearing or speech impairments

FTB 4192 PIT PC (REV 10-2012)

Figure 3.15: Top 500 website, landing page (10/6/2020)

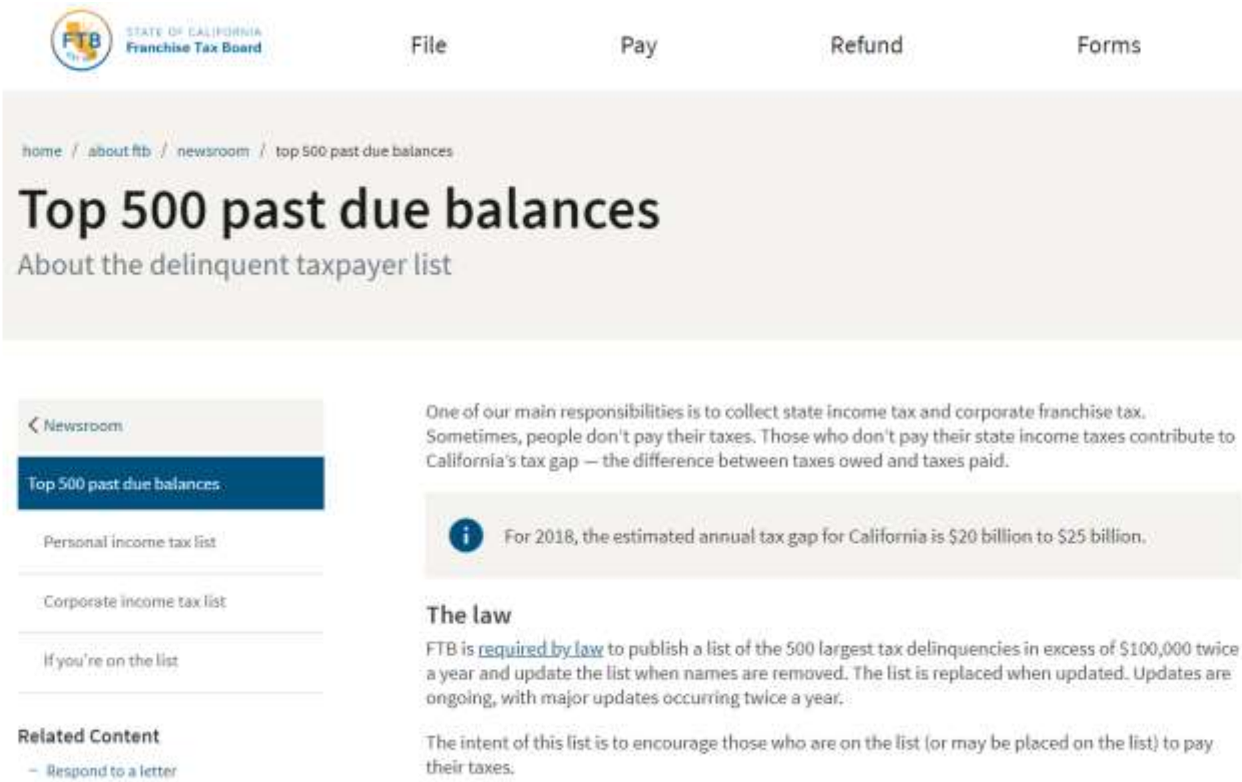
STATE OF CALIFORNIA Franchise Tax Board

File Pay Refund Forms

home / about ftb / newsroom / top 500 past due balances

Top 500 past due balances

About the delinquent taxpayer list

< Newsroom

Top 500 past due balances

Personal income tax list

Corporate income tax list

If you're on the list

Related Content

- Respond to a letter

One of our main responsibilities is to collect state income tax and corporate franchise tax. Sometimes, people don't pay their taxes. Those who don't pay their state income taxes contribute to California's tax gap — the difference between taxes owed and taxes paid.

For 2018, the estimated annual tax gap for California is $20 billion to $25 billion.

The law

FTB is required by law to publish a list of the 500 largest tax delinquencies in excess of $100,000 twice a year and update the list when names are removed. The list is replaced when updated. Updates are ongoing, with major updates occurring twice a year.

The intent of this list is to encourage those who are on the list (or may be placed on the list) to pay their taxes.

Figure 3.16: Top 500 list website, top balances (10/6/2020)


This is a list of the 500 largest tax delinquencies over $100,000. By law, we must publish this list at least twice a year.

Last updated: 10/06/2020

Name	Address	Subtotal	Total	Lien	License	Status	Number
Moreland, Peggy J & Terry L	Bakersfield, CA 93306	$5,306,836.86	$5,306,836.86	01/24/2008	Contractor's State License Board	Expired	856954
					Contractor's State License Board	Expired	362166
Cooksey, Jimmy D	Bowling Green, KY 42104	$2,403,194.62	$2,403,194.62	04/25/2008			
Amin, Joseph & Sharona	Beverly Hills, CA 90210	$1,730,698.65	$1,730,698.65	04/14/2014	Board of Pharmacy	Active	0034252
Patrick, William L & Susan K	Cody, WY 82414	$1,648,546.31	$1,648,546.31	05/31/2019			

3.10. Appendix B: Additional Figures and Tables

Figure 3.17: Behavior before and after the official letter, recipients vs. non-recipients (including the full range of balances)

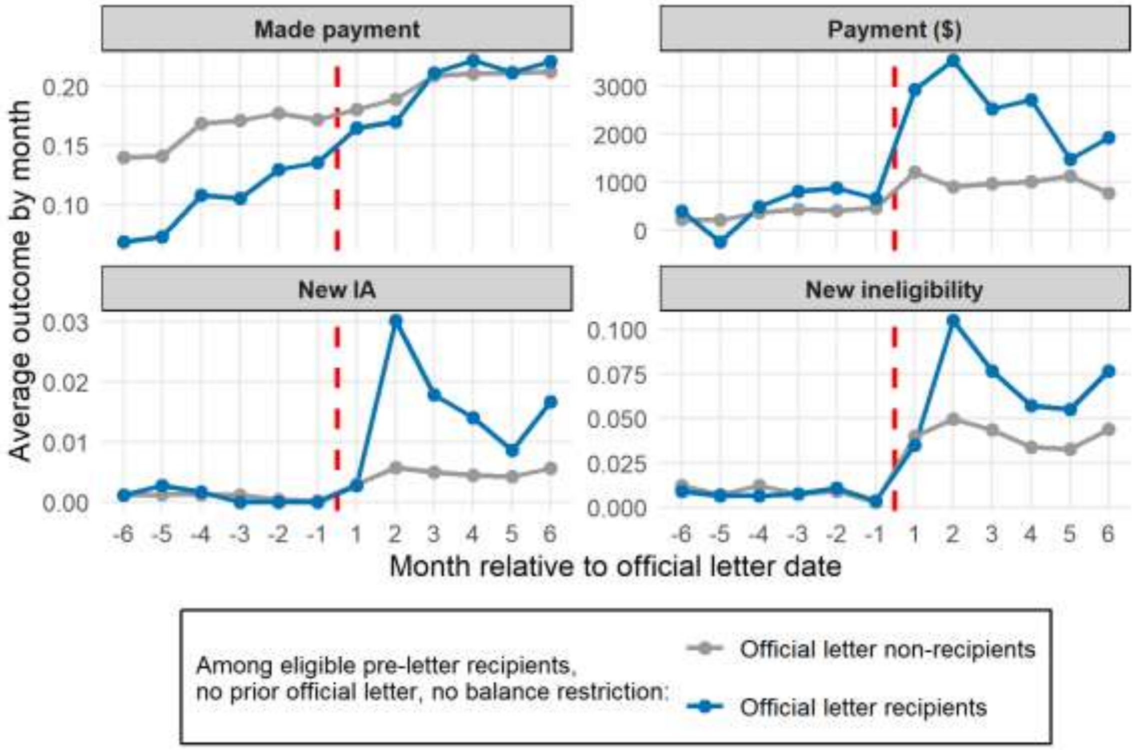

Notes: This figure compares the behavior of two groups of taxpayers around the date of the official letter. In gray are taxpayers who are eligible for publication and received a pre-letter, but did not receive an official letter. In blue are similar individuals (eligible pre-letter recipients) who did receive an official letter. In this figure we exclude individuals who previously received an official letter (i.e., we focus on first-time recipients), and we do not restrict to only those within the cutoff range.

Figure 3.18: Behavior before and after the official letter, testing importance of data filters

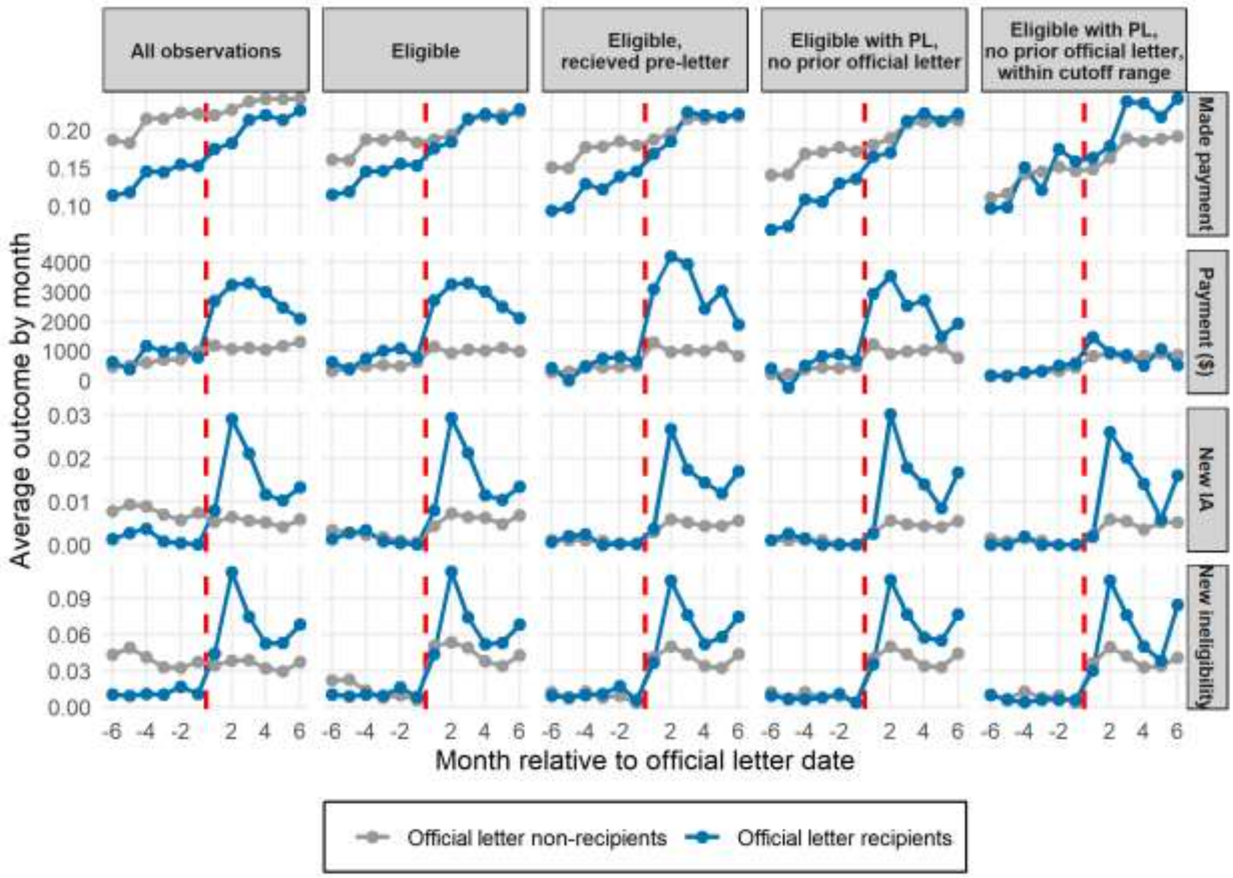

Notes: This figure compares the behavior of two groups of taxpayers around the date of the official letter, when applying various data filters. Going from left to right, the first column includes all observations in the dataset; the second restricts to those eligible for publication (based on their most recent status codes); the third further restricts to those receiving a pre-letter; the fourth further restricts to those who received no prior official letter; and the fifth adds the final restriction that balance fall between \$150,000 and \$230,000, the range of cutoff values inside of which treatment is quasi-random.

Table 3.5: *Official letter event study results*

	Dependent variables:			
	Made payment	Payment amount ($)	New IA	New ineligibility
Treatment dummies for each month relative to official letter date				
-6	-0.0587***	-432.341***	-0.0026***	-0.0128***
	(0.0147)	(80.7791)	(0.0003)	(0.0045)
-5	-0.0567***	-447.2311***	-0.0026***	-0.0168***
	(0.0146)	(65.0554)	(0.0003)	(0.0035)
-4	-0.0045	-311.7061***	-0.0006	-0.0188***
	(0.0169)	(92.2779)	(0.0020)	(0.0029)
-3	-0.0346**	-263.1644	-0.0026***	-0.0168***
	(0.0157)	(251.9457)	(0.0003)	(0.0036)
-2	0.0196	-88.8724	-0.0026***	-0.0168***
	(0.0179)	(207.8576)	(0.0003)	(0.0035)
-1	0.0035	-9.7269	-0.0026***	-0.0168***
	(0.0173)	(375.4987)	(0.0003)	(0.0035)
1	0.0075	873.7468	-0.0006	0.0073
	(0.0175)	(620.2508)	(0.0020)	(0.0077)
2	0.0236	359.4303	0.0235***	0.0816***
	(0.0182)	(381.6091)	(0.0071)	(0.0137)
3	0.0818***	275.1936	0.0175***	0.0535***
	(0.0201)	(368.0232)	(0.0063)	(0.0119)
4	0.0798***	-88.3425	0.0114**	0.0274***
	(0.0199)	(163.1266)	(0.0053)	(0.0098)
5	0.0618***	477.5197	0.0034	0.0153*
	(0.0195)	(464.5481)	(0.0035)	(0.0086)
6	0.0858***	-65.9806	0.0134**	0.0615***
	(0.0201)	(109.2293)	(0.0056)	(0.0125)
Balance	-0.0002	1.1828	0.0000	0*
($ thousands)	(0.0002)	(1.4382)	(0.0000)	(0.0000)
April publicatior	-0.0139***	-100.5314*	-0.0006	0.0021**
	(0.0038)	(56.5107)	(0.0004)	(0.0010)
Intercept	0.196***	404.4991	0.0054***	0.0307***
	(0.0425)	(259.2547)	(0.0018)	(0.0048)
Observations	75,696	75,696	75,696	75,696
R2	0.0021	0.0003	0.0026	0.0044
Mean dep var.	0.1576	569.4367	0.0031	0.0246

Notes: *p<0.1; **p<0.05; ***p<0.01. Standard errors clustered by taxpayer are reported in parentheses.

Figure 3.19: Estimates of treatment by AGI quintiles

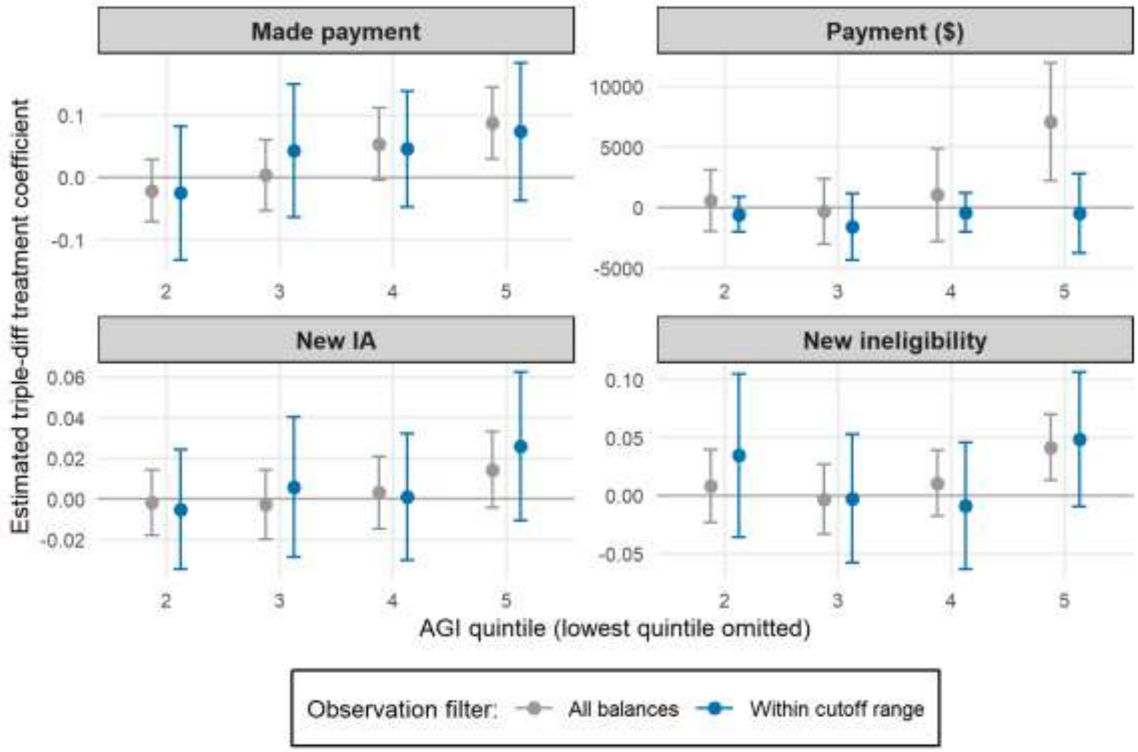

Notes: This figure presents the coefficient estimates for separate triple difference models testing AGI quintiles (i.e., the estimate for the coefficient on treatment X post X quintile). The lowest quintile is omitted. 95% confidence intervals are shown around the point estimates. In blue are the estimates using only observations within the cutoff range. These can be compared to the estimates in gray, from models using the full range of balance observations.

Table 3.6: Official letter triple difference results, heterogeneity tests (within cutoff range)

	Dependent variables:			
	Made payment	Payment amount ($)	New IA	New ineligibility
On-time filer	**0.0074**	**-125.51**	**0.017*** **	**0.0353*** **
(1/0)	(0.0230)	600.10	(0.0061)	(0.0130)
Among those with filed prior-year returns:				
Has business income	**0.0134**	**1864.61** **	**0.0187**	**0.0262**
(1/0)	(0.0382)	914.62	(0.0122)	(0.0202)
AGI	**0.0151**	**834.08**	**0.0066** **	**0.0109**
($ millions)	(0.0217)	1260.68	(0.0030)	(0.0154)
Above median AGI	**0.0693** **	**-759.18**	**0.0224** **	**0.0281***
(1/0)	(0.0304)	875.86	(0.0098)	(0.0150)
Non-negative AGI	**-0.0016**	**-725.57**	**0.0018**	**-0.0079**
(1/0)	(0.0380)	1100.72	(0.0151)	(0.0270)

Notes: This table summarize regression results for regressions in which the outcome variable appears in the column headers across the top. Reported coefficients are for a triple-interaction term between post, treatment, and the variable listed in the row labels in the leftmost column. Only those with balance within the range of cutoffs (roughly $150,000 to $230,000) are included. An "on-time filer" is one who filed a tax return the year prior to receiving treatment. AGI = annual gross income; the median in our sample is roughly $40,000. Standard errors clustered by taxpayer. *p<0.1; **: p<.05; ***: p <.01.

Table 3.7: Official letter triple difference results, heterogeneity tests (full sample)

	Dependent variables:			
	Made payment	Payment amount ($)	New IA	New ineligibility
On-time filer	**0.0130**	**852.36**	**0.0157*** **	**0.0168** **
(1/0)	(0.0122)	940.00	(0.0033)	(0.0070)
Among those with filed prior-year returns:				
Has business income	**-0.0204**	**2410.45***	**0.0099**	**0.0027**
(1/0)	(0.0202)	1432.27	(0.0063)	(0.0103)
AGI	**0.0055**	**772.23**	**0.0012**	**0.0033**
($ millions)	(0.0048)	495.89	(0.0022)	(0.0042)
Above median AGI	**0.0606*** **	**3053.57** **	**0.018*** **	**0.0266*** **
(1/0)	(0.0168)	1318.25	(0.0052)	(0.0083)
Non-negative AGI	**0.0320**	**3508.82** **	**0.0034**	**0.0216***
(1/0)	(0.0232)	1752.72	(0.0078)	(0.0121)

Notes: This table summarize regression results for regressions in which the outcome variable appears in the column headers across the top. Reported coefficients are for a triple-interaction term between post, treatment, and the variable listed in the row labels in the leftmost column. An "on-time filer" is one who filed a tax return the year prior to receiving treatment. AGI = annual gross income; the median in our sample is roughly $40,000. Standard errors clustered by taxpayer. *p<0.1; **: p<.05; ***: p <.01.

Figure 3.20: Official letter long-term effect estimates (comparing control approaches)

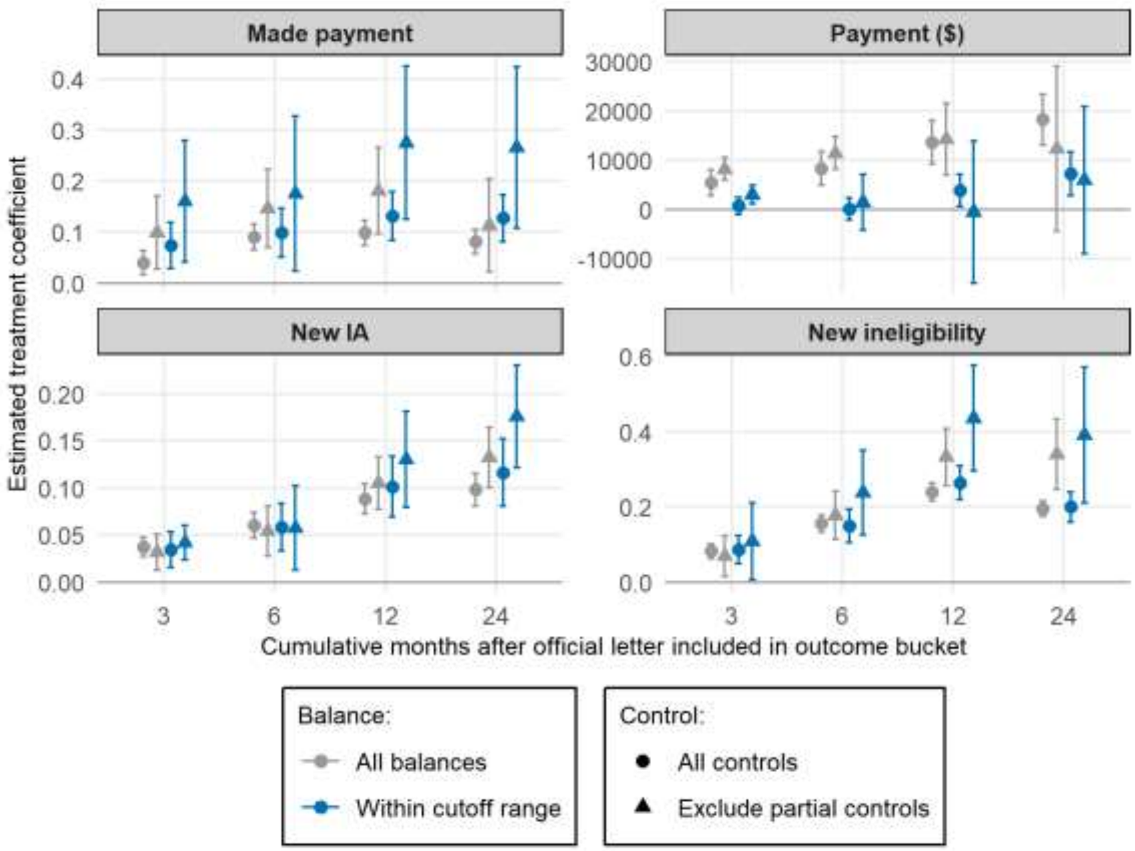

Notes: This figure summarizes estimated coefficients for regressions in which the variable in the grey bar is the outcome, cumulatively defined over the number of months post-official letter noted on the X-axis. Blue markers are for regressions in which we restrict the sample to observations with balances falling in between the historic lowest and highest Top 500 cutoff balances. Estimates shown with circles are the same as those presented above in Figure 3.9; estimates shown with triangles exclude those deemed "partial controls" (individuals in the control group who receive a letter during the subsequent two years).

Table 3.8: Official letter long-term payment effects (full range of observations)

	Dependent variables: Cumulative payment amount post-official-letter ($)			
	3 months	6 months	12 months	24 months
Official letter	**5388.73*** (1,318.31)**	**8332.44*** (1,746.36)**	**13608.8*** (2,234.60)**	**18283.3*** (2,628.68)**
Balance ($ thousands)	1.13 (0.74)	1.85 (1.15)	2.61* (1.50)	2.63* (1.59)
April publication	-2048*** (510.48)	-2487.13*** (684.55)	-2406.15*** (740.21)	-2034.9** (831.59)
Intercept	3869.5*** (405.42)	6838.05*** (562.87)	10802.19*** (703.61)	15299*** (835.56)
Observations	21,074	21,074	21,074	21,074
R2	0.0039	0.0049	0.0070	0.0075
Mean dep var.	3,626	6,835	11,519	16,620

Notes: *p<0.1; **p<0.05; ***p<0.01. Standard errors clustered by taxpayer are shown in parentheses. Outcomes measured as cumulative payments between time of treatment and time following, as listed in column headers. The full range of observations are included.

Figure 3.21: Behavior around the publication date, among official letter recipients

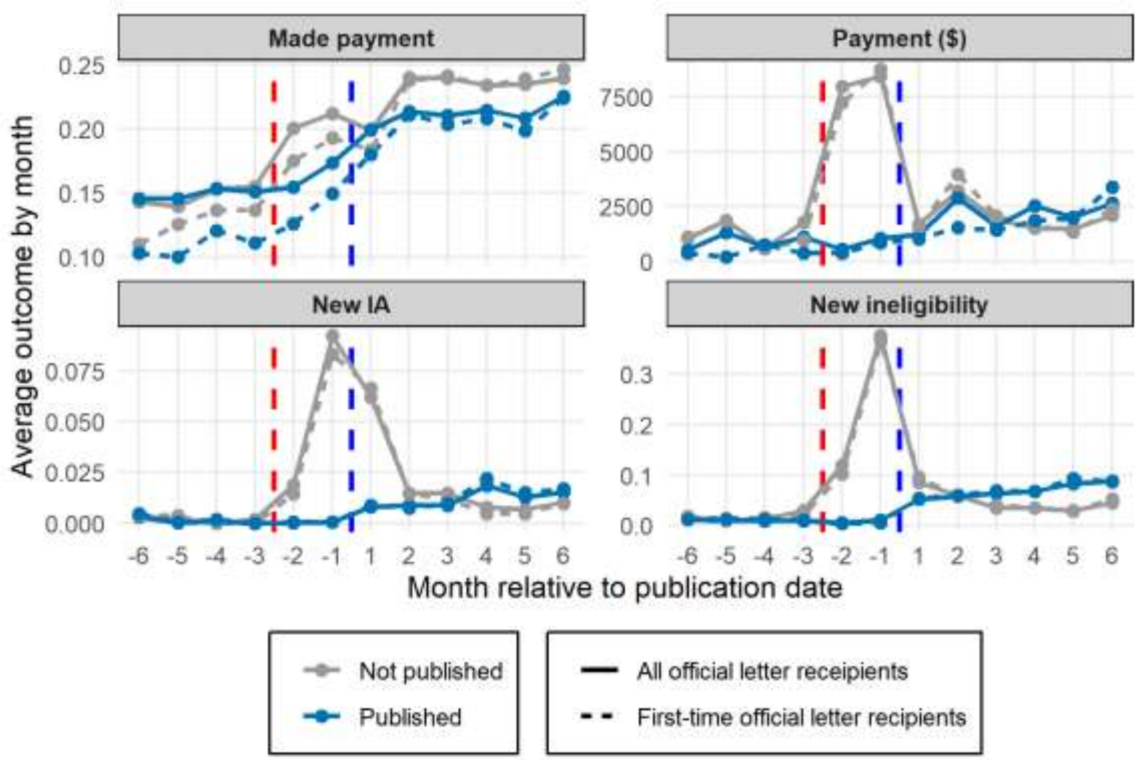

Notes: This figure compares the publication time series when including all official letter recipients (solid lines) and just first-time recipients (dashed lines). The patterns are similar.

3.11. Appendix C: Robustness Analyses

In this section, we test whether the findings discussed above are sensitive to the assumptions and data filters that define our main specification. Beginning with the pooled diff-in-diff approach, we test the sensitivity of the results to the following changes:

- Removing the requirement for same-cycle pre-letter receipt;
- Expanding the ineligibility definition to include any ineligible status during the prior month;
- Removing the eligibility filter entirely;
- Dropping any non-recipients above the cutoff value;
- Including individuals who received official letters for prior cycles;
- Narrowing the range of included values to those within the 2nd highest and lowest cycle cutoff values
- Not controlling for balance
- Including relative month as a control
- Considering different pre/post time windows

Table 3.9: *Official letter difference-in-difference, data filter robustness checks*

	Dependent variables:			
	Made payment	Payment amount ($)	New IA	New ineligibility
Main specification	**0.0223***	**155.77**	**0.0119*****	**0.0283*****
	(0.0118)	305.11	(0.0032)	(0.0066)
Robustness to data filters:				
Pre-letter: Don't require same-cycle pre-letter	0.0262**	105.67	0.0111***	0.03***
	(0.0105)	251.18	(0.0031)	(0.0062)
Eligibility: Exclude those with any past-month ineligible status change	0.0251	-98.75	0.0079**	0.0159**
	(0.0153)	316.79	(0.0036)	(0.0075)
Eligibility: Don't apply any eligible status filter	0.0161	221.16	0.0181***	0.0611***
	(0.0117)	303.78	(0.0033)	(0.0067)
Eligibility: Drop any non-recipients above that-cycle cutoff value	0.0216*	75.15	0.0119***	0.0279***
	(0.0119)	311.62	(0.0033)	(0.0066)
Prior Top500 experience: Don't restrict to first-time recipients	0.0165	319.36	0.0119***	0.0273***
	(0.0116)	315.10	(0.0030)	(0.0063)
Balance range: Use narrower range of included balances	0.0193	21.75	0.0138***	0.0307***
	(0.0135)	321.73	(0.0038)	(0.0075)
Balance control: Don't control for balance due	0.0223*	155.77	0.0119***	0.0283***
	(0.0118)	305.11	(0.0032)	(0.0066)
Time trend control: Include relative month covariate	-0.0479*	585.08	-0.0037	-0.0164
	(0.0271)	1010.72	(0.0057)	(0.0158)

Notes: *p<0.1; **p<0.05; ***p<0.01. Standard errors clustered by taxpayer are shown in parentheses. This table reports the coefficient on the treat*post term from each alternative specification of the difference-in-difference approach. See Table 3.2 and the accompanying text above for a discussion of the full model.

Table 3.10: Official letter difference-in-difference, time window robustness checks

	Dependent variables:			
	Made payment	Payment amount ($)	New IA	New ineligibility
Main specification	**0.0223***	**155.77**	**0.0119*****	**0.0283*****
	(0.0118)	305.11	(0.0032)	(0.0066)
Robustness to observation window:				
Restrict to one month pre/post	0.0021	516.19	-0.0001	-0.0090
	(0.0170)	742.71	(0.0021)	(0.0088)
Restrict to two months pre/post	-0.0027	190.21	0.0101***	0.0246***
	(0.0135)	442.42	(0.0038)	(0.0084)
Expand to four months pre/post	0.0273**	28.56	0.0115***	0.0277***
	(0.0116)	233.87	(0.0028)	(0.0055)
Expand to five months pre/post	0.0312***	58.24	0.0095***	0.0232***
	(0.0112)	211.83	(0.0023)	(0.0046)
Expand to six months pre/post	0.0367***	-6.27	0.01***	0.0267***
	(0.0111)	181.18	(0.0021)	(0.0043)

Notes: *p<0.1; **p<0.05; ***p<0.01. Standard errors clustered by taxpayer are shown in parentheses. This table reports the coefficient on the treat*post term from each alternative specification of the difference-in-difference approach. See Table 3.2 and the accompanying text above for a discussion of the full model.

Bibliography

Ahn, Julie Jee Eun. 2015. "The HEART Act of 2008: To Deter Expatriation, But Why More Expatriation." *North Carolina Journal of International Law and Commercial Regulation* 4: 1021-84.

Akcigit, Ufuk, Salomé Baslandze, and Stefanie Stantcheva. 2016. "Taxation and the International Mobility of Inventors." *American Economic Review* 106 (10): 2930-81.

Allingham, Michael G., and Agnar Sandmo. 1972. "Income Tax Evasion: A Theoretical Analysis." *Journal of Public Economics* 1: 323-338.

Andreoni, James. 1992. "IRS as Loan Shark: Tax Compliance with Borrowing Constraints." *Journal of Public Economics* 49 (1): 35-46.

Angaretis, Chad, Brian Galle, Paul R. Organ, and Allen Prohofsky. 2021. "Does Shaming Pay? Evaluating California's Top 500 Tax Delinquent Publication Program." *Working paper.*

Angrist, Joshua. 2005. "Instrumental Variables in Experimental Criminological Research: What, Why, and How?" *NBER Technical Working Paper 314.*

Arsenault, Steven J. 2009. "Surviving a Heart Attack: Expatriation and the Tax Policy Implications of the New Exit Tax." *Akron Tax Journal* 24 (2).

Belnap, Andrew, Jacob Thornock, and Braden Williams. 2021. "Hidden Wealth and Automatic Information Sharing." *Working paper.*

Benoit, David. 2012. "Facebook Co-Founder Saverin Fires Back At 'Misinformation'." *The Wall Street Journal*, May 17.

Bhargava, Saurabh, and Dayanand Manoli. 2015. "Psychological Frictions and the Incomplete Take-Up of Social Benefits: Evidence from an IRS Field Experiment." *American Economic Review* 105: 3489-3529.

Blank, Joshua D. 2014. "Collateral Compliance." *University of Pennsylvania Law Review* 162 (4): 719-800.

Bott, Kristina M., Alexander W. Cappelen, Erik O. Sorenson, and Bertil Tungodden. 2019. "You’ve Got Mail: A Randomized Field Experiment on Tax Evasion." *Management Science* 66: 2801-2819.

Bottan, Nicolas L., and Ricardo Perez-Truglia. 2020. "Betting on the House: Subjective Expectations and Market Choices." *NBER Working Paper 27412*.

Bottan, Nicolas L., and Ricardo Perez-Truglia. 2020. "Betting on the House: Subjective Expectations and Market Choices." *NBER Working Paper No. 27412*.

Chetty, Raj. 2009. "Is the Taxable Income Elasticity Sufficient to Calculate Deadweight Loss? The Implications of Evasion and Avoidance." *American Economic Journal: Economic Policy* 1 (2): 31-52.

Christians, Allison. 2017. "Buying In: Residence and Citizenship by Investment." *Saint Louis University Law Journal* 62: 51-72.

Collins, Brett, Alan Plumley, Ishani Roy, Alex Turk, Terry Ashley, and Jeff Wilson. 2018. "The Effectiveness of the Notice of Federal Tax Lien and Alternative IRS Letters in Resolving Individual Tax Debt." *IRS Bulletin.*

Craig, Bradford. 2012. "Congress, Have a Heart: Practical Solutions to Punitive Measures Plaguing the HEART Act’s Expatriate Inheritance Tax." *Temple International & Comparative Law Journal* 1: 69-102.

Cranor, Taylor, Jacob Goldin, and Tatiana Homonoff. 2020. "Communicating Tax Penalties to Delinquent Taxpayers: Evidence from a Field Experiment." *National Tax Journal* 73: 331-360.

Dahlby, Bev. 2008. *The Marginal Cost of Public Funds: Theory and Application.* Cambridge: MIT Press.

De Simone, Lisa, Rebecca Lester, and Kevin Markle. 2020. "Transparency and Tax Evasion: Evidence from the Foreign Account Tax Compliance Act (FATCA)." *Journal of Accounting Research* 58 (1): 105-53.

Dharmapala, Dhammika. 2016. "Cross-border tax evasion under a unilateral FATCA regime." *Journal of Public Economics* 141: 29-37.

Dorrenberg, Philipp, and Jan Schmitz. 2017. "Tax Compliance and Information Provision – A Field Experiment with Small Firms." *Journal of Behavioral Economics for Policy* 1: 47-54.

Dwenger, Nadja, and Lukas Treber. 2019. "Shaming for Tax Enforcement: Evidence from a New Policy." *CEPR Discussion Paper DP13194.*

Feldstein, Martin. 1999. "Tax Avoidance and the Deadweight Loss of the Income Tax." *Review of Economics and Statistics* 81 (4): 674-680.

Florida, Richard. 2011. "America's Great Passport Divide." *The Atlantic*, March 15.

Freeman, Richard B. 2006. "People Flows in Globalization." *Journal of Economic Perspectives* 20 (2): 145-70.

Galle, Brian. 2013. "Tax, Command...Or Nudge? Evaluating the New Regulation." *Texas Law Review* 92: 837-894.

Galle, Brian, and Murat Mungan. 2021. "Optimal Enforcement with Heterogeneous Private Costs of Punishment." *American Law & Economics Review* (in press).

Gemmell, Norman, and Marisa Ratto. 2018. "The Effects of Penalty Information on Compliance: Evidence from a New Zealand Field Experiment." *National Tax Journal* 71: 548-588.

Gibson, John, and David McKenzie. 2011. "Eight Questions about Brain Drain." *Journal of Economic Perspectives* 25 (3): 107-28.

Goldin, Jacob, and Daniel Reck. 2020. "Revealed Preference Analysis with Framing Effects." *Journal of Political Economy* 128: 2759-2795.

Goldin, Jacob, Tatiana Homonoff, Rizwan Javaid, and Brenda Schafer. 2021. "Tax Filing and Take-Up: Experimental Evidence on Tax Preparation Outreach and EITC Particiaption." *SIEPR Working Paper No. 21-007.*

Gonçalves, Pedro. 2018. "IRS to Revoke 260,000 American Passports." *International Investment*, October 16.

Hallsworth, Michael, John A. List, Robert D. Metcalfe, and Ivo Vlaev. 2017. "The Behavioralist as Tax Collector: Using Natural Field Experiments to Enhance Tax Compliance." *Journal of Public Economics* 148: 14-31.

Hartner, Martina, Silvia Rechberger, Erich Kirchler, and Alfred Schabmann. 2008. "Procedural Fairness and Tax Compliance." *Economic Analysis of Policy* 38: 137-152.

Heckman, James J., Seong Hyeok Moon, Rodrigo Pinto, Peter A. Savelyev, and Adam Yavitz. 2010. "The Rate of Return to the HighScope Perry Preschool Program." *Journal of Public Economics* 94 (1-2): 114-128.

Hendren, Nathaniel, and Ben Sprung-Keyser. 2020. "A Unified Welfare Analysis of Government Policies." *The Quarterly Journal of Economics* 135 (3): 1209-1318.

Holz, Justin E., John A. List, Alejandro Zentner, Marvin Cardoza, and Joaquin Zentner. 2020. "The $100 Million Nudge: Increasing Tax Compliance of Businesses and the Self-Employed Using a Natural Field Experiment." *NBER Working Paper No. 27666.*

Iyer, Govind S., Phillip M.J. Reckers, and Debra L. Sanders. 2010. "Increasing Tax Compliance in Washington State: A Field Experiment." *National Tax Journal* 63: 7-32.

Johannesen, Niels, Patrick Langetieg, Daniel Reck, Max Risch, and Joel Slemrod. 2020. "Taxing Hidden Wealth: The Consequences of U.S. Enforcement Initiatives on Evasive Foreign Accounts." *American Economic Journal: Economic Policy* 12 (3): 312-46.

Joint Committee on Taxation. 1995. "Issues presented by proposals to modify the tax treatment of expatriation: A report." JCS-17-95.

Kaplow, Louis. 2008. *The Theory of Taxation and Public Economics.* Princeton: Princeton University Press.

Keen, Michael, and Joel Slemrod. 2017. "Optimal Tax Administration." *Journal of Public Economics* 152: 133-142.

Kenchington, David, and Roger White. 2021. "Income Tax Noncompliance and Professional License Suspension: Evidence from Missouri." *Working paper.*

Kim, Young Ran (Christine). 2017. "Considering 'Citizenship Taxation': In Defense of FATCA." *Florida Tax Review* 20 (5): 335-70.

Kleven, Henrik, Camille Landais, and Emmanuel Saez. 2013. "Taxation and International Migration of Superstars: Evidence from the European Football Market." *American Economic Review* 103 (5): 1892-1924.

Kleven, Henrik, Camille Landais, Mathilde Muñoz, and Stefanie Stantcheva. 2020. "Taxation and Migration: Evidence and Policy Implications." *Journal of Economic Perspectives* 34 (2): 119-42.

Kuchumova, Yulia. 2018. "A Collateral Tax Sanction: When Does it Mimic a Welfare-Improving Tag?" *International Tax and Public Finance* 25 (4): 841-866.

Kuchumova, Yulia. 2021. "Tax Debt Collection Enforcement: When Does Suspension of a Driver's License Help?" *Working paper.*

Kudrle, Robert T. 2015. "Expatriation: A Last Refuge for the Wealthy?" *Global Policy* 6 (4): 408-17.

Kwong, Yu Hang Sonny. 2009. "Catch Me If You Can: Relinquishing Citizenship for Taxation Purposes After the 2008 HEART Act." *Houston Business and Tax Journal* 9: 411-44.

Luttmer, Erzo M.P., and Monica Singhal. 2014. "Tax Morale." *Journal of Economic Perspectives* 28: 149-168.

Manolakas, Christine, and William L. Dentino. 2012. "The Exit Tax: A Move in the Right Direction." *William & Mary Business Law Review* 3: 341-418.

Mason, Ruth. 2016. "Citizenship Taxation." *Southern California Law Review* 89: 169-240.

Meiselman, Ben S. 2018. "Ghostbusting in Detroit: Evidence on Nonfilers from a Controlled Field Experiment." *Journal of Public Economics* 158: 180-193.

Mirrlees, James A. 1982. "Migration and Optimal Income Taxes." *Journal of Public Economics* 18 (3): 319-41.

Murphy, Kristina. 2003. "Procedural Justice and Tax Compliance." *Australian Journal of Social Issues* 38 (3): 379–408.

Nathan, Brad C., Ricardo Perez-Truglia, and Alejandro Zentner. 2021. "My Taxes are Too Darn High: Why do Houeholds Protest Their Taxes?" *NBER Working Paper 27816.*

Nathan, Brad C., Ricardo Perez-Truglia, and Alejandro Zentner. 2021. "My Taxes Are Too Darn High: Why Do Households Protest Their Taxes?" *NBER Working Paper No. 27816.*

Organ, Paul R., Alex Ruda, Joel Slemrod, and Alex Turk. 2021. "Do Collateral Sanctions Work? Evidence from the IRS' Passport Certification and Revocation Process." *NBER Working Paper No. 29029.*

Perez-Truglia, Ricardo, and Guillermo Cruces. 2017. "Partisan Interactions: Evidence from a Field Experiment in the United States." *Journal of Political Economy* 125 (4): 1208-43.

Perez-Truglia, Ricardo, and Guillermo Cruces. 2017. "Partisan Interactions: Evidence from a Field Experiment in the United States." *Journal of Political Economy* 125 (4): 1208-1243.

Perez-Truglia, Ricardo, and Ugo Troiano. 2018. "Shaming Tax Delinquents." *Journal of Public Economics* 167: 120-137.

Phillips Erb, Kelly. 2017. "If You Owe Taxes, You Could Lose Your Passport And Your Ability To Fly." *Forbes*, January 26.

Polinsky, A. Mitchell, and Steven Shavell. 2000. "The Economic Theory of Public Enforcement of Law." *Journal of Economic Literature* 38 (1): 45-76.

Reichenberg Sherr, Eileen. 2008. "New Expatriation Tax Rules Enacted." *The Tax Adviser*, August.

Richards, Paul J. 2016. "Delays, costs mount for Canadians renouncing U.S. citizenship." *The Globe and Mail*, February 9.

Rigg, Stephanie Cellini, Fernando Ferreira, and Jesse Rothstein. 2010. "The Value of School Facility Investments: Evidence from a Dynamic Regression Discontinuity Design." *The Quarterly Journal of Economics* 125: 215-261.

Sheppard, Lee. 2009. "Don't Ask, Don't Tell, Part 4: Ineffectual Information Sharing." *Tax Notes*, March 23.

Slemrod, Joel. 2019. "Tax Compliance and Enforcement." *Journal of Economic Literature* 57 (4): 904-954.

Slemrod, Joel, and Christian Gillitzer. 2013. *Tax Systems.* Cambridge: MIT Press.

Slemrod, Joel, and Shlomo Yitzhaki. 1987. "The Optimal Size of a Tax Collection Agency." *Scandinavian Journal of Economics* 89 (2): 183-192.

Slemrod, Joel, Obeid Ur Rehman, and Mazhar Waseem. 2020. "How Do Taxpayers Respond to Public Disclosure and Social Recognition Programs? Evidence from Pakistan." *The Review of Economics and Statistics* (in press).

Treasury Inspector General for Tax Administration (TIGTA). 2019. "Implementation of the Passport Provisions of the FAST Act Was Generally Successful, and the Internal Revenue Service Is Working on Objective Criteria for Passport Revocations." *Reference Number: 2019-30-068,* September 19.

Turk, Alex, John Iuranich, Stacy Orlett, and Saurabh Datta. 2016. "Resolving Unpaid Taxes and the Notice of Federal Tax Lien: Evidence from the Fresh Start Initiative." *IRS Bulletin.*

Tyler, T. R. 1999. "Why People Cooperate with Organizations: An Identity-Based Perspective." *Research in Organizational Behavior* 21: 201-246.

U.S. Government Accountability Office (GAO). 2011. "Federal Tax Collection: Potential for Using Passport Issuance to Increase Collection of Unpaid Taxes." *GAO-11-272.*

U.S. House of Representatives. 2015. "Conference Report on H.R. 22, Surface Transportation Reauthorization and Reform Act of 2015." https://www.congress.gov/congressional-record/2015/12/01/house-section/article/H8679-2.

Urist, Jacoba. 2012. "No Taxes, No Travel: Why the IRS Wants the Right to Seize Your Passport." *The Atlantic*, April 17.

Westin, Richard A. 2000. "Expatriation and Return: An Examination of Tax-Driven Expatriation by United States Citizens, and Reform Proposals." *Virginia Tax Review* 20 (1): 75-190.